Optimization and Reoptimization in Scheduling Problems

arXiv:1509.01630v1 [cs.DS] 4 Sep 2015

Yael Mordechai

Optimization and Reoptimization in Scheduling Problems

Research Thesis

Submitted in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Computer Science

Yael Mordechai

Submitted to the Senate of the
Technion - Israel Institute of Technology

Tamuz 5775 Haifa June 2015

This research was carried out under the supervision of Prof. Hadas Shachnai, in the Computer Science Department.

The generous financial help of the Technion is gratefully acknowledged.

Contents

Abstract 1

1 Introduction 2

1.1 Scheduling on Parallel Machines . . . 2

1.1.1 Scheduling Models . . . 2

1.1.2 Related Work . . . 3

1.2 Fixed Parameter Algorithms . . . 5

1.2.1 Parametrized Scheduling Problems . . . 6

1.2.2 Related Work . . . 6

1.3 Combinatorial Reoptimization . . . 7

1.3.1 Reoptimization in Scheduling . . . 7

1.3.2 Related Work . . . 7

1.4 Main Results . . . 8

2 Makespan Minimization for Fully-Feasible Instances 10

2.1 Preliminaries . . . 10

2.2 Approximation Algorithm for Fully-Feasible Instances . . . 11

2.2.1 Approximation Algorithm . . . 12

2.3 A Better Bound for the Restricted Assignment Problem . . . 18

2.3.1 Approximation Algorithm . . . 19

3 Fixed-Parameter Algorithms for Scheduling on Unrelated Machines 22

3.1 Preliminaries . . . 22

3.2 Parametrized Approximation Scheme for Scheduling on Unrelated Machines 23

3.3 An FPT Algorithm for Graph-Balancing . . . 24

4 Reoptimization Algorithms for Scheduling Problems 28

4.1 Preliminaries . . . 28

4.2 A $(1, 1+\epsilon)$-Reapproximation Algorithm for Makespan Minimization on Identical Machines . . . 29

4.2.1 Algorithm . . . 29

4.2.2 Analysis . . . 30

4.3 A $(1, 1+\epsilon)$-Reapproximation Algorithm for Makespan Minimization on Uniform Machines . . . 35

4.3.1 Algorithm . . . 35
4.3.2 Analysis . . . 37

5 Conclusions and Future Work **40**
5.1 Summary . . . 40
5.2 Future Work . . . 40

List of Figures

2.1 Example of the construction in the rounding 12
2.2 The bipartite graph G_σ. By changing the orientation of the path $i_3 \to j_5 \to i_5$, we remove job j_5 from i_5 and schedule it on machine i_3. 19

List of Tables

1.1 Known results for makespan minimization on unrelated machines. 5
1.2 Known FPT algorithms for scheduling. 7
1.3 Our contribution for reoptimization in scheduling. 9

Abstract

Parallel machine scheduling has been extensively studied in the past decades, with applications ranging from production planning to job processing in large computing clusters. In this work we study some of these fundamental optimization problems, as well as their parameterized and reoptimization variants.

We first present improved bounds for job scheduling on unrelated parallel machines, with the objective of minimizing the latest completion time (or, makespan) of the schedule. We consider the subclass of *fully-feasible* instances, in which the processing time of each job, on any machine, does not exceed the minimum makespan. The problem is known to be hard to approximate within factor 4/3 already in this subclass. Although fully-feasible instances are hard to identify, we give a polynomial time algorithm that yields for such instances a schedule whose makespan is better than twice the optimal, the best known ratio for general instances. Moreover, we show that our result is robust under small violations of feasibility constraints.

We further study the power of parameterization. In a parameterized optimization problem, each input comes with a fixed parameter. Some problems can be solved by algorithms (or approximation algorithms) that are exponential only in the size of the parameter, while polynomial in the input size. The problem is then called *fixed parameter tractable (FPT)*, since it can be solved efficiently (by an FPT algorithm or approximation algorithm) for constant parameter values. We show that makespan minimization on unrelated machines admits a *parameterized approximation scheme*, where the parameter used is the number of processing times that are large relative to the latest completion time of the schedule. We also present an FPT algorithm for the graph-balancing problem, which corresponds to the instances of the *restricted assignment* problem where each job can be processed on at most 2 machines.

Finally, motivated by practical scenarios, we initiate the study of *reoptimization* in job scheduling on identical and uniform machines, with the objective of minimizing the makespan. We develop *reapproximation* algorithms that yield in both models the best possible approximation ratio of $(1+\epsilon)$, for any $\epsilon > 0$, with respect to the minimum makespan.

Chapter 1

Introduction

1.1 Scheduling on Parallel Machines

Consider the following fundamental problem in scheduling theory. We are given a set $\mathcal{J}$ of n independent jobs that must be scheduled without preemption on a collection $\mathcal{M}$ of m parallel machines. If job j is scheduled on machine i, the processing time required is p_{ij}, which is a positive integer, for every $i \in \mathcal{M}$ and $j \in \mathcal{J}$. The total time used by machine $i \in \mathcal{M}$, or the *load* on machine i, is the sum of the processing times for the jobs assigned to i, and the *makespan* of an assignment is the maximum load over all the machines. The objective is then to find a schedule, which assigns each job to exactly one machine, such that the makespan is minimized.

1.1.1 Scheduling Models

The wide literature on scheduling often distinguishes between the following scheduling models.

Identical Machines. Job processing times are identical across the machines, i.e., $p_{ij} = p_j$ for all $j \in \mathcal{J}$ and $i \in \mathcal{M}$.

Uniform Machines. Each machine i has a speed s_i. The length of job j on machine i is some uniform processing time p_j scaled by the speed s_i, i.e., $p_{ij} = \frac{p_j}{s_i}$ for all $j \in J$ and $i \in \mathcal{M}$.

Unrelated Machines: Each job j may have an arbitrary processing time $p_{ij} \geq 0$ on machine i, for $j \in \mathcal{J}$ and $i \in \mathcal{M}$.

While makespan minimization is known to be NP-hard in all models (even for $m = 2$) [LK79], the first two models are considered somewhat easier, since the problem can be approximated efficiently in both up to some ϵ factor, for any $\epsilon > 0$ (see Section 1.1.2). In contrast, in the unrelated machines model, the problem becomes hard to approximate within a factor better than $\frac{3}{2}$. Moreover, since 1990, when the state of the art 2-approximation algorithm was presented by Lenstra, Shmoys and Tardos [LST90], there was no significant improvement on either the upper or lower bound, although the problem was consistently investigated. This led researchers to consider special cases and improving the bound of 2, either by a constant factor, or by some function of the input parameters (see review in Section 1.1.2).

In this work, we consider the subclass of *fully-feasible* instances. We say that an instance is fully-feasible if job processing times, p_{ij}, do not exceed the length of the optimal schedule for the instance, for every job $j \in \mathcal{J}$ and machines $i \in \mathcal{M}$. Observe that an optimal schedule never assigns a job to machine on which its length is greater than the makespan of an optimal schedule; thus, from an optimal scheduler's viewpoint, if p_{ij} exceeds the optimal makespan then p_{ij} is considered to be ∞.

We also consider instances that are almost fully-feasible, that is, for any job j, the number of machines on which job j is not feasible (i.e., has processing time larger than the length of the optimal makespan) is relatively small.

When considering real-life applications, the general model of unrelated machines, which makes no assumptions on job processing times, seems too broad. Indeed, such applications usually deal with fully-feasible (or almost fully-feasible) workloads, as they commonly handle relatively large sets of jobs.

Let T_{opt} and L_{opt} denote the optimal makespan and the minimal average machine load over optimal assignments, respectively. For heterogeneous workloads of a huge number of jobs, in which the makespan is counted in months or even years, the processing time of a given job is negligible compared to the makespan; for such workloads, we have $p_{ij} \ll T_{opt}$. In this case, an algorithm of [ST93], yields a schedule of makespan at most $T_{opt} + p_{max}$, where $p_{max} = max_{ij} p_{ij}$, is more suitable. However, for smaller sets of jobs, p_{ij} can be large relative to T_{opt}, such that $L_{opt} < p_{max}$, and for these instances our algorithm is the state of the art. Relevant applications for such workloads are e.g., job packing in warehouse-scale [VKW14], large-scale clustering [VP+15] and applications in parallel design patterns such as Fork-Join and MapReduce (see, e.g., [DMN12, LL+12]).

1.1.2 Related Work

Identical and Uniform Machines The problem of makespan minimization on identical or uniform machines is known to be NP-hard [GJ79]. A *polynomial-time approximation scheme (PTAS)* is a family of algorithms $\{A_\epsilon : \epsilon > 0\}$, where A_ϵ is a $(1+\epsilon)$-approximation algorithm that runs in time polynomial in the input size but is allowed to be exponential in $\frac{1}{\epsilon}$. An *efficient polynomial-time approximation scheme (EPTAS)* is a PTAS with running time $f(\frac{1}{\epsilon}) poly(|\mathcal{I}|)$, where $|\mathcal{I}|$ is the input size, (for some function f), while a *fully polynomial-time approximation scheme (FPTAS)* runs in time $poly(\frac{1}{\epsilon}, |\mathcal{I}|)$. Since the scheduling problem is NP-hard in the strong sense already on identical machines (as it contains bin packing and 3-Partition as special cases) [GJ79], we cannot hope for an FPTAS. For identical machines, Hochbaum [H96] and Alon et al. [AA+98] gave an EPTAS with running time $f(\frac{1}{\epsilon}) + O(n)$, where f is doubly exponential in $\frac{1}{\epsilon}$, and for uniform machines, Jansen [J10] gave an EPTAS with running time $2^{O(1/\epsilon^2 log(1/\epsilon)^3)} + poly(|\mathcal{I}|)$.

Unrelated Machines A classic result in scheduling theory is the Lenstra-Shmoys-Tardos 2-approximation algorithm for makespan minimization [LST90]. They also proved that the problem is NP-hard to approximate within a factor better than $\frac{3}{2}$. Gairing et al. [GMW07] presented a more efficient, combinatorial 2-approximation algorithm based on flow techniques. Shchepin and Vakhania [SV05] showed that the rounding technique used in [LST90] can be modified to derive an improved ratio of $2 - \frac{1}{m}$. Shmoys and Tardos [ST93] showed an approximation algorithm that yields a schedule of makespan at most the length of an optimal schedule plus the largest processing time of any job in the instance.

Although makespan minimization on unrelated machines is a major open problem in scheduling theory, and is extensively studied, there was no significant progress on either the upper or lower bound for over two decades, since the publication of [LST90]. This led researchers to consider interesting special cases and improving the upper bound for them. A well known special case is the *restricted assignment problem*, where jobs have processing times $p_{ij} \in \{p_j, \infty\}$. Svensson [S12] gave a polynomial-time algorithm that approximates the optimal makespan of the restricted assignment problem within a factor of $\frac{33}{17} + \epsilon \approx 1.94 + \epsilon$ for $\epsilon > 0$, and also presented a local search algorithm that will eventually find a schedule of the mentioned approximation guarantee, but is not known to converge in polynomial-time. Gairing et al. [GL+04] presented a combinatorial $(2 - \frac{1}{p_{max}})$-approximation algorithm for the restricted assignment problem.

Ebenlendr et al. considered in [EKS08] the *graph balancing* problem, a special case of the restricted assignment problem where each job j has a finite processing time, $p_{ij} < \infty$, on at most two machines. The paper gives an elaborate 1.75-approximation algorithm for the problem. The authors also show that the problem is hard to approximate within a factor less than $\frac{3}{2}$ even on bounded degree graphs, i.e., when the maximum degree is some constant. In the *unrelated graph balancing* problem, introduced by Versache and Weiss [VW14], each job can be assigned to at most two machines, but processing times are not restricted. They showed that this subclass of instances constitutes the core difficulty for the *linear programming* formulation of makespan minimization on unrelated machines, often used as a first step in obtaining approximate solutions. Specifically, they showed that the strongest known LP-formulation, namely, the configuration-LP, has an integrality gap of 2.

Vakhania et al. [VMH14] considered makespan minimization on unrelated machines for the subclass of instances where job lengths can take only two values, p and q, which are fixed positive integers, such that $p < q$. They presented a polynomial-time algorithm that uses linear programming with absolute approximation factor of q (i.e., all schedules have makespan at most $OPT + q$). Page [P14] considered restricted assignment instances with processing times in a fixed interval, $[p, q]$, and gave a $\frac{q}{p}$-approximation algorithm, and a $\frac{3}{2}$-approximation algorithm for the case where $p_{ij} \in \{1, 2, 3\}$. Chakrabarty et al. [CK+15] considered instances with two types of jobs: *long* and *short*, namely, $p_{ij} \in \{1, \epsilon\}$ for some $\epsilon > 0$. They obtained a $(2 - \delta)$-approximation algorithm for such instances.

Shmoys and Tardos [ST93] considered the *generalized assignment problem (GAP)*, where each job j incurs a cost of $c_{ij} > 0$ when assigned on machine i, and the objective is to minimize the makespan and the total cost. The paper [ST93] presents a polynomial-time algorithm that finds a schedule of makespan at most twice the optimum with optimal cost.

A summary of the known results for unrelated machines is given in Table 1.1.

Result	Authors	Restrictions on the Unrelated Model
2-approximation	[LST90]	
Hard for factor $< \frac{3}{2}$	[LST90]	
A bound of $T_{opt} + p_{max}$	[ST93]	
$(2 - \frac{1}{m})$-approximation	[SV05]	
Integrality gap ≤ 1.95	[S12]	$p_{ij} \in \{p_j, \infty\}$
1.75-approximation	[EKS08]	$p_{ij} \in \{p_j, \infty\}$, $p_{ij} < \infty$ on < 2 machines
$(2 - \delta)$-approximation, $\delta > 0$	[CK+15]	$p_{ij} \in \{1, \epsilon\}$ for some $\epsilon > 0$
A bound of $T_{opt} + q$	[VMH14]	$p_{ij} \in \{p, q\}$ for some $p < q$
$\frac{q}{p}$ & $\frac{3}{2}$-approximation	[P14]	$p_{ij} \in [p, q]$ & $p_{ij} \in \{1, 2, 3\}$
A bound of $T_{opt} + L_{opt}$	This Work	$p_{ij} \leq T_{opt}$
A bound of $T_{opt} + \frac{L_{opt}}{\varphi}$	This Work	$\varphi \geq \frac{L}{T}$, for minimal feasible values of T & L
A bound of $p_{max} + \frac{L_{opt}}{\varphi}$	This Work	$p_{ij} \in \{p_j, \infty\}$

Table 1.1: Known results for makespan minimization on unrelated machines.

1.2 Fixed Parameter Algorithms

Parameterized complexity is a branch of computational complexity theory that focuses on classifying computational problems according to their inherent difficulty with respect to multiple parameters of the input. The complexity of a problem is then measured as a function in those parameters. This allows the classification of NP-hard problems on a finer scale than in the classical setting, where the complexity of a problem is only measured by the number of bits in the input (see, e.g., [DF12]).

Under the assumption that $P \neq NP$, there exist many natural problems that require super-polynomial running time when complexity is measured in terms of the input size only, but that are computable in a time that is polynomial in the input size and exponential (or worse) in a parameter k. Hence, if k is fixed at a small value, and the growth of the function over k is relatively small, then such problems can still be considered “tractable” despite their traditional classification as “intractable”.

Some problems can be solved exactly, or approximately, by algorithms that are exponential only in the size of a fixed parameter while polynomial in the size of the input. Such an (approximation) algorithm is called a *fixed-parameter tractable (FPT)* (approximation) algorithm, because the problem can be (approximately) solved efficiently for small values of the fixed parameter.

Problems in which some parameter k is fixed are called parameterized problems. A parameterized problem that allows for such an FPT algorithm is said to be a *fixed-parameter* *tractable* problem and belongs to the class FPT.

1.2.1 Parametrized Scheduling Problems

Despite the fundamental nature of scheduling problems, and the clear advantages of fixed-parameter algorithms, no such algorithms are known for many of the classical scheduling problems. One obstacle towards obtaining positive results appears to be that, in contrast to most problems known to be fixed-parameter tractable, scheduling problems involve many numerical input data (e.g., job processing times, release dates, job costs), which alone causes many problems to be NP-hard, thus ruling out fixed-parameter algorithms.

1.2.2 Related Work

While minimizing the makespan on identical and uniform machines admits an EPTAS, see e.g., [AA+98, J10], the running times of these approximation schemes usually have a bad dependence on ϵ. In addition, these problems are strongly NP-hard [GJ79], therefore we cannot hope to obtain an FPTAS. These considerations call for finding which scheduling problems are fixed-parameter tractable (FPT). This amounts to identifying instance-dependent parameters k that allow for algorithms that find optimal solutions in time $f(k) \cdot poly(|\mathcal{I}|)$, for instances $\mathcal{I}$ and some function f depending only on k.

As for scheduling on unrelated machines, an FPTAS is known [HS76], but only when assuming a fixed number of machines. Note that, if the number of machines or the number of processing times are constant, the problem is still NP-hard [LST90], and thus no FPT-algorithms can exist for these choices of parameters. This motivates us to identify instance-dependent parameters k, while assuming an arbitrary number of machines, that allow for better FPT approximation algorithms.

Marx [M11] proposed the research direction of scheduling with rejection, where the parameter k is given with the input for the scheduling problem, and the solution has to schedule all but k jobs. This direction was explored recently by Mnich and Wiese [MW13], who presented for the first time a fixed-parameter algorithms for classical scheduling problems such as makespan minimization, scheduling with job-dependent cost functions and scheduling with rejection. For the problem of makespan minimization on identical machines, the paper [MW13] presents an FPT algorithm, where the parameter k defines an upper bound on the number of distinct processing times appearing in an instance. For the more general model of unrelated machines, the paper gives an FPT algorithm, using the number of machines and the number of distinct processing times as parameters.

Table 1.2 summarizes the known results.

Result	Parameters	Authors	Model
FPT algorithm	$max_j p_j$	[MW13]	identical machines
FPT algorithm	m, #distinct p_{ij}	[MW13]	unrelated machines
FPT approximation scheme	$\|\{(i,j) : p_{ij} > \epsilon T\}\|$, feasible T	This Work	unrelated machines
FPT algorithm	treewidth, degree	This Work	graph balancing

Table 1.2: Known FPT algorithms for scheduling.

We are not aware of any other work in this area.

1.3 Combinatorial Reoptimization

Reoptimization problems naturally arise in many real-life scenarios. Indeed, planned or unanticipated changes occur over time in almost any system. It is then required to respond to these changes quickly and efficiently. Ideally, the response should maintain high performance while affecting only a small portion of the system. Thus, throughout the continuous operation of such a system, it is required to compute solutions for new problem instances, derived from previous instances. Since the transition from one solution to another incurs some cost, a natural goal is to have the solution for the new instance close to the original one (under certain distance measure).

We use the reoptimization model developed by Shachnai et al. [STT12]. In this model, we say that $\mathcal{A}$ is an (r, ρ)-reapproximation algorithm if it achieves a ρ-approximation for the optimization problem, while incurring a transition cost that is at most r times the minimum cost required for solving the problem optimally.

1.3.1 Reoptimization in Scheduling

We consider instances of scheduling problems that can change dynamically over time. Our goal is to compute assignments within some guaranteed approximation for the new problem instances, derived from the previous instances. Since the transition from one assignment to another incurs some cost (for example, the cost of pausing the execution of a job on one machine and resuming its execution on another), an additional goal is to have the solution for the new instance *close* to the original one (under a certain distance measure).

1.3.2 Related Work

Reoptimization Shachnai et al. [STT12] presented reapproximation algorithms for several non-trivial classes of optimization problems. This includes a *fully polynomial time reapproximation schemes (FPTRS)* for DP-benevolent problems, reapproximation algorithms for metric Facility Location problems, and $(1, 1)$-reoptimization algorithm for polynomially solvable subset-selection problems.

Junosza-Szaniawski et al. [JLR15] considered a variant of a recoloring problem, called the r-Color-Fixing. They investigated the problem of finding a proper r-coloring of a graph, which is "most similar" to some given initial solution, i.e. the number of vertices that have to be recolored is minimum possible. They provide a $(1, 1)$-reoptimization algorithm for the problem. More work on reoptimization can be found e.g., in [J15, AE+14].

Reoptimization in Scheduling Baram and Tamir [BT14] considered the problem of scheduling on identical machines with the objective of minimizing the total flow time, i.e., minimizing $\sum_{j\in\mathcal{J}} C_j$, where C_j is the completion time of job j. They presented an algorithm that yields an optimal solution using the minimal possible transition cost, and an algorithm that outputs the best possible schedule, using a given limited budget for the transition, for several classes of instances.

Bender et al. [BF+13] focused on a scheduling problem where each job is unit-sized and has a time window in which it can be executed. Jobs are dynamically added and removed from the system. They presented an algorithm that reschedules only $O(min\{log^* n, log^* \Delta\})$ jobs for each job that is inserted or deleted from the system, where n is the number of active jobs, and Δ is the size of the largest window.

1.4 Main Results

Our first contribution is in the study of makespan minimization on unrelated machines, which leads to approximation algorithms with performance guarantees better that 2, the best known bound for general instances. In particular, for the subclass of fully-feasible instances, we present (in Chapter 2) an LP-based algorithm that achieves makespan at most $T_{opt} + L_{opt}$, where T_{opt} and L_{opt} are the minimal makespan and minimal average machine load for this makespan, respectively. This result is better than twice the minimal makespan for instances naturally arising in real-life applications. It also improves the minimal makespan plus the maximum processing time of a job, for instances where the average machine load is smaller than the maximum processing time of any job. We show that our algorithm is *robust* in the sense that it achieves an improved makespan also for instances that are *almost* fully-feasible. In such instances, each job may exceed the minimal makespan on a *small* number of machines. Formally, we define the *feasibility parameter* of a general instance $\mathcal{I}$, denoted $\varphi(\mathcal{I})$, as the minimal fraction of machines on which a job has a processing time at most T_{opt}, i.e., $\varphi(\mathcal{I}) = \frac{min_{j\in\mathcal{J}}|\{i\in\mathcal{M}:p_{ij}\leq T_{opt}\}|}{m}$. We present an algorithm that yields, for instances with large *enough* feasibility parameter, a schedule of makespan at most $T_{opt} + \frac{L_{opt}}{\varphi(\mathcal{I})}$. For instances of the restricted assignment problem, i.e., for instances with processing times $p_{ij} \in \{p_j, \infty\}$, we show that a bound of $p_{max} + \frac{L_{opt}}{\varphi(\mathcal{I})}$ is obtained by an efficient and simple combinatorial algorithm, where p_{max} is the largest processing time of any job in the instance.

We further study the power of parameterization and present an FPT approximation scheme, i.e., a $(1+\epsilon)$-FPT approximation algorithm, for makespan minimization on unrelated machines parametrized by the number of machine-job pairs, $(i,j) \in \mathcal{M} \times \mathcal{J}$, such that $p_{ij} > \epsilon T$, for some makespan candidate T. We also show that the graph-balancing problem, parameterized by treewidth and the maximum degree of the graph, is in FPT. These results are presented in Chapter 3.

Our third contribution is reapproximation algorithms for the reoptimization variants of makespan minimization on identical and uniform machines, which are studied here for the first time. Specifically, we develop $(1, 1+\epsilon)$-reapproximation algorithms, namely, algorithms that achieve a ratio of $(1+\epsilon)$ to the minimum makespan, and the minimum transition cost, in both the identical machines and the uniform machines models, where transition costs can take values in $\{0,1\}$. For the uniform case, we assume that the ratio between the highest and the lowest machine speeds is bounded by some constant. Thus, our algorithms achieve the best possible ratio with respect to the makespan objective in these models. For the unrelated machines model, we note that an algorithm of Shmoys and Tardos [ST93] can be used to obtain a $(1,2)$-reapproximation, thus matching the best known bound for makespan minimization also in this model. We summarize the results for reoptimization in scheduling in Table 1.3. The results are given in Chapter 4.

Result	costs	Authors	Model
$(1,2)$-reapproximation algorithm	arbitrary	[ST93]	unrelated machines
$(1,1+\epsilon)$-reapproximation algorithm	{0,1}	This Work	identical machines
$(1,1+\epsilon)$-reapproximation algorithm	{0,1}	This Work	uniform machines with $\frac{s_1}{s_m} \leq b$

Table 1.3: Our contribution for reoptimization in scheduling.

Chapter 2

Makespan Minimization for Fully-Feasible Instances

2.1 Preliminaries

Consider a scheduling instance $\mathcal{I} = (\mathcal{J}, \mathcal{M})$, consisting of a set of m machines $\mathcal{M}$ and a set of n jobs $\mathcal{J}$ with non-negative integers p_{ij} denoting the processing time of job $j \in \mathcal{J}$ on machine $i \in \mathcal{M}$. An *assignment* of the jobs to the machines is a bijection $\sigma : \mathcal{J} \to \mathcal{M}$ where $\sigma(j) = i$ if and only if job j is assigned to machine i. For any assignment σ, the *load* on machine i under assignment σ, denoted as $load_\sigma(i)$, is the sum of processing times for the jobs that were assigned to machine i. Thus, $load_\sigma(i) = \sum_{j \in \mathcal{J}:\alpha(j)=i} p_{ij}$. The *makespan* of an assignment σ, denoted by $T(\sigma)$, is the maximum load over all the machines. Thus, $T(\sigma) = max_{i \in \mathcal{M}} load_\sigma(i)$. The *average machine load* of an assignment σ, denoted by $L(\sigma)$, is given by $L(\sigma) = \frac{\sum_{i \in \mathcal{M}} load_\sigma(i)}{m}$. Given an instance $\mathcal{I} = (\mathcal{J}, \mathcal{M})$, we denote by T_{opt} the optimal makespan, i.e., $T_{opt} = min_{\sigma:\mathcal{M} \to \mathcal{J}} T(\sigma)$, and we denote by L_{opt} the minimum average machine load for an optimal assignment, i.e., $L_{opt} = min_{\sigma^*:\mathcal{M} \to \mathcal{J}, T(\sigma^*)=T_{opt}} L(\sigma^*)$.

Given an instance $\mathcal{I} = (\mathcal{J}, \mathcal{M})$, we say that a job j is *feasible* on machine i, if and only if $p_{ij} \leq T_{opt}$. The *feasibility parameter* of $\mathcal{I}$, denoted $\varphi(\mathcal{I})$, is the minimal fraction of feasible machines for any given job, i.e., $\varphi(\mathcal{I}) = min_{j \in \mathcal{J}} \frac{|\{i \in \mathcal{M} : j \text{ is feasible on } i\}|}{m}$. Thus, every job $j \in \mathcal{J}$ is feasible on at least $\varphi(\mathcal{I}) \cdot m$ machines. In this terms, an instance $\mathcal{I}$ is *fully-feasible* if and only if $\varphi(\mathcal{I}) = 1$. We often omit $\mathcal{I}$ in the notation if it is clear from the context, and refer to the feasibility parameter as φ.

Given an instance $\mathcal{I} = (\mathcal{J}, \mathcal{M})$, an assignment $\sigma : \mathcal{J} \to \mathcal{M}$, two positive integers L and T and some real number $\gamma \geq 1$, we denote by $Bad(\sigma, \gamma) \subseteq \mathcal{M}$ the subset of machines i with $load_\sigma(i) > T + \gamma \cdot L$, by $Good(\sigma, \gamma) \subseteq \mathcal{M}$ the subset of machines i with $load_\sigma(i) \leq \gamma \cdot L$, and for all $j \in \mathcal{J}$, by $Good_j(\sigma, \gamma) \subseteq Good(\sigma, \gamma)$ the set of machines from $Good(\sigma, \gamma)$ that are also legal for job j. For every $i \in \mathcal{M}$ we denote by $j^i_{max}(\sigma) = argmax \{p_{ij} : \sigma(j) = i\}$ the job with the longest processing time, assigned, by σ, on machine i.

2.2 Approximation Algorithm for Fully-Feasible Instances

Given positive integers L and T, let x_{ij} be an indicator to the assignment of job j on machine i. Consider the following linear program.

$$
\begin{aligned}
LP(T,L): \qquad & \frac{1}{m}\sum_{i=1}^{m}\sum_{j=1}^{n} p_{ij}x_{ij} \leq L \\
& \sum_{i=1}^{m} x_{ij} = 1, && \text{for } j = 1, ..., n \\
& \sum_{j=1}^{n} p_{ij}x_{ij} \leq T, && \text{for } i = 1, ..., m \\
& x_{ij} \geq 0, && \text{for } i = 1, ...m\,, j = 1, ..., n
\end{aligned}
$$

One can see that integer solutions to the above LP are in one to one correspondence with assignments $\sigma : \mathcal{M} \to \mathcal{J}$ of average machine load at most L and makespan at most T.

Theorem 2.1. *If $LP(T,L)$ is feasible for some $L \leq T$, then there is a polynomial-time algorithm that yields a schedule σ with*

1. $L(\sigma) \leq L$
2. $load_\sigma(i) \leq T + max\{p_{ij} : \sigma(j) = i\}$, $\forall i \in \mathcal{M}$

Proof. Let $\bar{x} = (x_{ij} : i \in \mathcal{M}, | \in \mathcal{J})$ be a fractional solution to $LP(T,L)$. We round $\bar{x}$ to an integer solution using the following classic rounding technique, also used in [ST93]. Let $k_i = \left\lceil \sum_{j\in\mathcal{J}} x_{ij} \right\rceil$. Each machine is partitioned into k_i sub-machines $v_{i,s}$, $s = 1, ..., k_i$.

An edge weighted, bipartite graph $B = (W, V; E)$ is constructed, where $W = \{w_j : j = 1, ..., n\}$ representing the jobs and $V = \{v_{is} : i = 1, .., m,\ s = 1, ..., k_i\}$ representing the sub-machines. The edges are constructed in the following way.

Consider the nodes v_{is} as bins of unit capacity and the nodes w_j, as pieces of size x_{ij}. For every machine $i \in \mathcal{M}$, consider the nodes in non-increasing order of the processing time of the corresponding job, on machine i. For convenience, assume $p_{i1} \geq p_{i2} \geq ... \geq p_{in}$. An edge (w_j, v_{is}) with cost p_{ij} is constructed if and only if a positive fraction of x_{ij} is packed in the bin v_{is}. The packing of the bins (and construction of the edges) is done in the following way. The bins $v_{i1}, ..., v_{ik_i}$ are packed one by one, with the pieces in the order $p_{i1}, p_{i2}, ..., p_{in}$. While v_{is} is not totally packed, (otherwise we consider the next bin) we continue packing the next piece such that if its size, x_{ij}, fits v_{is} (without causing an overflow) it is packed to v_{is}, else, if it causes an overflow, only a fraction $\beta > 0$ of x_{ij} is packed to v_{is}, consuming all the remaining volume of v_{is}, and the remaining part of $(1-\beta)x_{ij}$ is packed in $v_{i,s+1}$. Figure 2.1 gives a pictorial example of this construction.

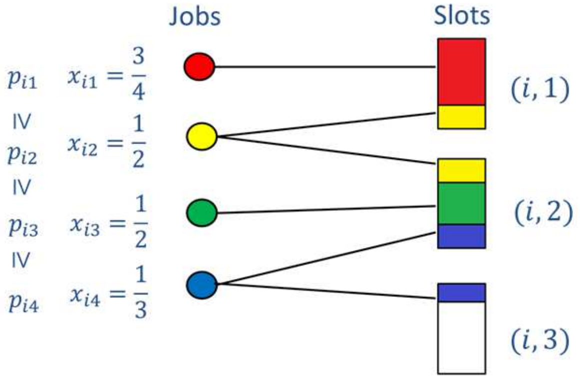
Jobs
Slots
p_{i1} $x_{i1} = \frac{3}{4}$
$\geq$
p_{i2} $x_{i2} = \frac{1}{2}$
$\geq$
p_{i3} $x_{i3} = \frac{1}{2}$
$\geq$
p_{i4} $x_{i4} = \frac{1}{3}$
$(i, 1)$
$(i, 2)$
$(i, 3)$

Figure 2.1: Example of the construction in the rounding

The rounding is done by finding a minimum-cost integer matching $M \in E$ that matches all job nodes, and for every edge $(w_j, v_{is}) \in M$, set $\sigma(j) = i$ i.e., assign job j on machine i.

By taking a minimum-cost integer matching it is guaranteed that $\frac{1}{m}\sum_{i=1}^{m}\sum_{j:\sigma(j)=i} p_{ij} \leq L$, or $L(\sigma) \leq L$. By the construction of the graph it is guaranteed that the load on machine i, for every $i \in \mathcal{M}$, is at most $max\{p_{ij} : \sigma(j) = i\} + \sum_{s=2}^{k_i} p_{is}^{max}$, where $p_{is}^{max} = max\{p_{ij} : (w_j, v_{is}) \in E\}$. Since $\sum_{s=2}^{k_i} p_{is}^{max} \leq T$ (for a detailed proof, see [ST93]), we get that $load_\sigma(i) \leq max\{p_{ij} : \sigma(j) = i\} + T$ for all $i \in \mathcal{M}$. $\square$

This result will be helpful in our approximation algorithm.

2.2.1 Approximation Algorithm

Recall that an instance $\mathcal{I}$ is *fully-feasible* if and only if $p_{ij} \leq T_{opt}$ for every job $j \in \mathcal{J}$ and machine $i \in \mathcal{M}$. Although fully-feasible instances are hard to identify, we give a polynomial-time algorithm that yields for such instances an assignment whose makespan is better than twice the optimal makespan, the best known ratio for general instances.

Our main result is as follows.

Theorem 2.2. *There is a polynomial-time algorithm that yields for fully-feasible instances an assignment of makespan at most* $T_{opt} + L_{opt}$.

Note that always $L_{opt} \leq T_{opt}$. $L_{opt} = T_{opt}$ occurs only in case where *every* optimal assignment is perfectly balanced (i.e., the load on all the machines equals T_{opt}). Thus, in the typical case we get $L_{opt} < T_{opt}$, which means that in the typical case our algorithm guaranteed a bound that is strictly better that twice the optimum.

The following theorem shows that the problem remains hard already when considering only the class of fully-feasible instances.

Theorem 2.3. *Makespan minimization on unrelated machines is hard to approximate within factor* $< \frac{4}{3}$ *, already in the class of fully-feasible instances.*

The proof of the hardness result follows from a reduction from 3-dimensional matching as in [LST90].

We prove a stronger result, that also shows that our result is robust under small violations of the feasibility constraints. Namely, we show that we can get better bounds also when our instance is not fully-feasible but has the property that every job can be non-feasible on some *small* fraction of the machines. For this, we will need some definitions.

Definition 2.1. *Job* $j \in \mathcal{J}$ *is feasible on machine* $i \in \mathcal{M}$ *if* $p_{ij} \leq T_{opt}$

Definition 2.2. *The feasibility parameter of instance* $\mathcal{I}$ *is defined as*

$$\varphi(\mathcal{I}) = min_{j \in \mathcal{J}} \frac{|\{i \in \mathcal{M} : j \text{ is feasible on } i\}|}{m}$$

In these terms, we have that every job $j \in \mathcal{J}$ is feasible on at least $\varphi(\mathcal{I}) \cdot m$ machines. Moreover, an instance $\mathcal{I}$ is fully-feasible if and only if $\varphi(\mathcal{I}) = 1$.

Given an instance $\mathcal{I}$, we can fix L and T to be the minimal values such that $LP(T, L)$ is feasible in polynomial-time in the size of the input $|\mathcal{I}|$, using binary search on some feasible regions for L and T. At the end of this operation, L and T will satisfy $L \leq T$, $L \leq L_{opt}$ and $T \leq T_{opt}$.

Theorem 2.4. *There is a polynomial-time algorithm that, given the values* $L \leq T$*, yields for instances* $\mathcal{I}$ *with feasibility parameter* $\varphi(\mathcal{I}) \geq \frac{L}{T}$*, an assignment of makespan at most* $T_{opt} + \frac{L_{opt}}{\varphi(\mathcal{I})}$

Corollary 2.5. *The statement of theorem 2.2 follows directly from 2.4, since the feasibility parameter of fully-feasible instances is* 1*.*

Given a general scheduling instance $\mathcal{I}$, its optimal makespan, T_{opt}, nor its feasibility parameter, $\varphi(\mathcal{I})$, cannot be computed in polynomial-time, unless $P = NP$. However, let ℓ be the number of distinct processing times p_{ij} of an instance $\mathcal{I}$ and w.l.o.g assume that $p_1 \leq p_2 \leq ... \leq p_\ell$ are the distinct processing times of $\mathcal{I}$. If we denote $p_{\ell+1} = \infty$, then we see that T_{opt} can belong to exactly one of the regions $[p_t, p_{t+1})$ for some $t \in \{1, ..., \ell\}$. If $T_{opt} \in [p_t, p_{t+1})$ for some $t \in \{1, ..., \ell\}$, then job $j \in \mathcal{J}$ is feasible on some machine $i \in \mathcal{M}$ if and only if $p_{ij} \leq p_t$. Therefore, $T_{opt} \in [p_t, p_{t+1})$ if and only if $\varphi(\mathcal{I}) = \varphi_t \triangleq min_{j \in \mathcal{J}} \frac{|\{i \in \mathcal{M} : p_{ij} \leq p_t\}|}{m}$.

We prove the theorem by describing an algorithm that admits the desired bound on the makespan, for instances with *large* enough feasibility parameter. The first step of the algorithm is to find the minimal T and L such that $LP(T, L)$ is feasible. Next, it obtains a fractional solution to $LP(T, L)$ and rounds it to obtain an initial assignment σ, such that $L(\sigma) \leq L$ and $load_\sigma(i) \leq T + max\{p_{ij} : \sigma(j) = i\}$, $\forall i \in \mathcal{M}$, as in Theorem 2.1. Then, for each guess of the optimal makespan region and the corresponding feasibility parameter, it tries to fix the initial assignment to achieve a new assignment of makespan at most $T_{opt} + \frac{L_{opt}}{\varphi}$ (and unless the feasibility parameter is too large, it will succeed for the right guess). This is done by balancing the initial assignment such as to reduce the load of the overloaded machines to meet the desired makespan.

We first prove the following lemmas.

Lemma 2.6. *Let σ be an assignment for some instance $\mathcal{I} = (\mathcal{J}, \mathcal{M})$ such that $L(\sigma) \le L$, let $T \ge L$ and let $\gamma \ge 1$. Denote $k = |Bad(\sigma, \gamma)|$. Then*

1. $k < \frac{m}{\gamma+1}$.
2. $|Good(\sigma, \gamma)| > \left(1 - \frac{1}{\gamma}\right) \cdot m + \frac{k}{\gamma} \cdot \frac{T}{L}$.

Proof. Each machine $i \in Bad(\sigma, \gamma)$ has load greater than $T + \gamma \cdot L$, therefore $\sum_{i \in \mathcal{M}} load_\sigma(i) > k \cdot (T + \gamma \cdot L)$.

1. Assume that $k \ge \frac{m}{\gamma+1}$, then

$$\begin{aligned} \textstyle\sum_{i \in M} load_\sigma(i) &> k(T + \gamma L) \\ &\ge \tfrac{m}{\gamma+1}(T + \gamma L) \\ &= \tfrac{m}{\gamma+1}(T + \gamma L) \\ &\ge \tfrac{m}{\gamma+1}(L + \gamma L) \\ &= m \cdot L \end{aligned}$$

   The last inequality follows from the fact that $T \ge L$. Hence, the average machine load is greater than L, a contradiction. It follows that $k < \frac{m}{\gamma+1}$.

2. Let $|Good(\sigma, \gamma)| = l$. Then, there are $m - k - l$ machines having loads greater than γL. Assume that $l \le \left(1 - \frac{1}{\gamma}\right) m + \frac{k}{\gamma} \cdot \frac{T}{L}$, then

$$\begin{aligned} \textstyle\sum_{i \in M} load_\sigma(i) &> k(T + \gamma \cdot L) + (m - l - k)\,\gamma L \\ &= kT + (m - l)\,\gamma L \\ &\ge kT + \left(m - \left(1 - \tfrac{1}{\gamma}\right) m + \tfrac{k}{\gamma} \cdot \tfrac{T}{L}\right) \gamma L \\ &\ge kT + \left(\tfrac{m}{\gamma} + \tfrac{k}{\gamma} \cdot \tfrac{T}{L}\right) \gamma L \\ &= kT + (mL + kT) \\ &\ge mL \end{aligned}$$

   Hence, the average machine load is greater than L, a contradiction. It follows that $|Good(\sigma, \gamma)| \ge \left(1 - \frac{1}{\gamma}\right) \cdot m + \frac{k}{\gamma} \cdot \frac{T}{L}$.

□

Lemma 2.7. *Let $\mathcal{I} = (\mathcal{J}, \mathcal{M})$ be an instance with feasibility parameter φ. Let σ be an assignment for $\mathcal{I}$ with average machine load L and let $T \ge L$. If $\varphi \ge \frac{L}{T}$ then for every subset $A \subseteq Bad(\sigma, \frac{1}{\varphi})$, $|N(A)| \ge |A|$, where $N(A)$ is the set of neighbors of A in $G_{\sigma, \frac{1}{\varphi}}$.*

Proof. Let $\left|Bad(\sigma, \frac{1}{\varphi})\right| = k$. Since the number of illegal machines for any job j is at most $(1-\varphi)m$, the number of good machines for job j is at least the number of good machines minus its illegal machines (the worst case where all illegal machines for job j are contained in $Good(\sigma, \frac{1}{\varphi})$). Together with Lemma 2.6 we have

$$\begin{aligned}\left|Good_j(\sigma, \tfrac{1}{\varphi})\right| &\geq \left|Good(\sigma, \tfrac{1}{\varphi})\right| - (1-\varphi)m \\ &> \left(1 - \tfrac{1}{\varphi}\right) \cdot m + \tfrac{k}{\left(\frac{1}{\varphi}\right)} \cdot \tfrac{T}{L} - (1-\varphi)m \\ &= \varphi \cdot k\tfrac{T}{L} \\ &\geq k\end{aligned}$$

The last inequality follows from the fact that $\varphi \geq \frac{L}{T}$. Now, let $A \subseteq Bad(\alpha, \frac{1}{\varphi})$. Then $|A| \leq \left|Bad(\sigma, \frac{1}{\varphi})\right| = k$.

Recall that the set of neighbors of A is the set of machines that are good for all the jobs j^i_{max}, $i \in A$, i.e., $N(A) = \cup_{i\in A} Good_{j^i_{max}}(\sigma, \frac{1}{\varphi}) \subseteq Good(\sigma, \frac{1}{\varphi})$. Obviously $|N(A)| = \left|\cup_{i\in A} Good_{j^i_{max}}(\sigma, \frac{1}{\varphi})\right| \geq \left|Good_{j^i_{max}}(\sigma, \frac{1}{\varphi})\right|$ for some $i \in A$. It follows from the above that $|N(A)| \geq k$.

Since $|A| \leq k$ we have that $|N(A)| \geq |A|$. ☐

By Hall's Theorem [H35], there exist a perfect matching in $G_{\sigma,\frac{1}{\varphi}}$ if and only if for every $A \subseteq Bad(\sigma, \frac{1}{\varphi})$, $|N(A)| \geq |A|$. Thus, we have

Corollary 2.8. *There exists a perfect matching in* $G_{\sigma,\frac{1}{\varphi}}$.

By the above discussion, we can modify the initial assignment σ, by finding a perfect matching in $G_{\sigma,\frac{1}{\varphi}}$ and then transferring jobs from bad machines to their matching good machines. We describe this formally in algorithm $\mathcal{A}_{UM}$.

Algorithm 1 $\mathcal{A}_{UM}$

```
1. Use binary search to find the minimal $T$, such that $LP(T,T)$ is feasible. Next, with that $T$ fixed, search for the minimal $L$ such that $LP(T,L)$ is feasible.
2. Solve the linear relaxation $LP(T,L)$.
3. Round the solution to obtain an integral assignment $\sigma$ using a rounding technique as given in Theorem 2.1.
4. For every $t = 1, ..., \ell$, where $\ell$ is the number of distinct processing times of $\mathcal{I}$, guess that $T_{opt} \in [p_t, p_{t+1})$ and that $\varphi(\mathcal{I}) = \varphi_t$.
    (a) If $\varphi_t < \frac{L}{T}$, continue.
    (b) Otherwise, construct the bipartite graph $G_{\sigma,\frac{1}{\varphi_t}}$ and find a perfect matching of size $\left|Bad(\sigma, \frac{1}{\varphi_t})\right|$, if one exists. If not, continue.
    (c) Obtain a resulting assignment $\sigma'$ from $\sigma$ by transferring the longest job, $j^i_{max}$ from each machine $i \in Bad(\sigma, \frac{1}{\varphi_t})$, to its matching machine $i' \in Good(\sigma, \frac{1}{\varphi_t})$.
5. Return the assignment $\sigma'$ with minimal makespan.
```

Proof of Theorem 2.4. We show that the assignment output by Algorithm $\mathcal{A}_{UM}$ satisfies the statement of the theorem. Consider a general instance $\mathcal{I}$. By performing binary search on a feasible bounded region of the optimal makespan we can find the minimal T for which $LP(T,T)$ is feasible, and then by performing binary search on a feasible bounded region of the optimal average machine load we can find the minimal L for which $LP(T,L)$ is feasible. These integers satisfy that $T \leq T_{opt}$, $L \leq L_{opt}$ and $L \leq T$.

By Theorem 2.1, since $LP(T,L)$ is feasible then Step 3 is guaranteed to generate an assignment σ of $L(\sigma) \leq L$ and $load_\sigma \leq T + max\{p_{ij} : \sigma(j) = i\}$, for all $i \in \mathcal{M}$.

Let $\varphi \geq \frac{L}{T}$ be the feasibility parameter of $\mathcal{I}$. Then for some $t = 1, ..., \ell$, where ℓ is the number of distinct processing times in $\mathcal{I}$, $\varphi_t = \varphi$. Then, by Corollary 2.8, there exists a perfect matching in $G_{\sigma,\frac{1}{\varphi_t}}$ and we will find it in Step $4(b)$ in $\mathcal{A}_{UM}$. Let $k = \left|Bad(\sigma, \frac{1}{\varphi_t})\right|$ and let $M = \{(i_{b_1}, i_{g_1}), ..., (i_{b_k}, i_{g_k})\}$ be a perfect matching in $G_{\sigma,\frac{1}{\varphi_t}}$.

For any machine $i = 1, ..., m$, $load_\sigma(i) \leq T + max\{p_{ij} : \sigma(j) = i\}$. Let j^i_{max} be the largest job processed by σ on machine i. Then, $p_{i,j^i_{max}} = max\{p_{ij} : \sigma(j) = i\}$, thus removing j^i_{max} from machine i guarantees that the new load of machine i will be at most T.

As for the good machines, if i is a good machine for job j then $p_{ij} \leq p_t$. Therefore, transferring j to i will increase the load of i by at most p_t which is at most T_{opt} (since $T_{opt} \in [p_t, p_{t+1})$). Since the load of a good machine is at most $\frac{L}{\varphi_t}$, we have that after such a job transfer the load will be at most $T_{opt} + \frac{L}{\varphi_t}$.

The load on the rest of the machines stays unchained, i.e., $load_\sigma(i) \leq T + \frac{L}{\varphi_t}$, for all $i \notin Bad(\sigma, \frac{1}{\varphi_t}) \cup Good(\sigma, \frac{1}{\varphi_t})$.

Thus, by performing the large-jobs transfers for all pairs $(i_{b_s}, i_{g_s}) \in M$, $s = 1, ..., k$, we obtain a new assignment σ', with $load_{\sigma'}(i) \leq max\{T, T_{opt} + \frac{L}{\varphi_t}, T + \frac{L}{\varphi_t}\}$ for all $i \in \mathcal{M}$, which is at most $T_{opt} + \frac{L_{opt}}{\varphi_t}$, since $L \leq L_{opt}$ and $T \leq T_{opt}$. $\square$

Theorem 2.9. *The complexity of* $\mathcal{A}_{UM}$ *is* $O((nm)^{\frac{7}{2}} \cdot log^2(\sum_{(i,j)\in\mathcal{M}\times\mathcal{J}} p_{ij}))$.

Proof. We will show that Step 1 in $\mathcal{A}_{\mathcal{UM}}$ is the bottle-neck of the algorithm. In this step we perform a binary search on the feasible regions of T and L while solving $LP(T, L)$. Since the feasible region for both T and L is $[0, \sum_{(i,j)\in\mathcal{M}\times\mathcal{J}} p_{ij}]$, we solve the LP $O(log(\sum_{(i,j)\in\mathcal{M}\times\mathcal{J}} p_{ij}))$ times. Solving the LP can be done in time $O((nm)^{\frac{7}{2}} \cdot log(\sum_{(i,j)\in\mathcal{M}\times\mathcal{J}} p_{ij}))$ [K84], so overall this operation runs in $O((nm)^{\frac{7}{2}} \cdot log^2(\sum_{(i,j)\in\mathcal{M}\times\mathcal{J}} p_{ij}))$.

The number of vertices in the bipartite graph $G_{\sigma,\frac{1}{\varphi_t}}$ equals to the number of bad and good machines, which is at most m. From 2.6, the number of bad machines is at most $\frac{m}{\frac{1}{\varphi}+1}$, thus the number of edges in the bipartite graph is at most $\sum_{i\in Bad(\sigma,\frac{1}{\varphi})} \varphi m = \frac{\varphi}{\frac{1}{\varphi}+1} \leq \frac{\varphi}{2} \leq \frac{1}{2}$ (the last two inequalities are due to $\varphi \leq 1$). Therefore finding a perfect matching in $G_{\sigma,\frac{1}{\varphi_t}}$ can be done in time $O(\sqrt{|V_{\sigma,\frac{1}{\varphi_t}}|} \cdot |E_{\sigma,\frac{1}{\varphi_t}}|) = O(\sqrt{m})$ [MV80]. In $\mathcal{A}_{UM}$, we find a perfect matching for every $t = 1, .., \ell$, where $\ell \leq n \cdot m$ is the number of distinct processing times. Thus, the complexity of Step 4 sums to $O(nm\sqrt{m})$.

Step 2 is done in $O((nm)^{\frac{7}{2}} \cdot log(\sum_{(i,j)\in\mathcal{M}\times\mathcal{J}} p_{ij}))$ time and Step 3 is done in $O(n^3(m+n)^3)$ [EK72].

Therefore, the overall complexity of the algorithm is $O((nm)^{\frac{7}{2}} \cdot log^2(\sum_{(i,j)\in\mathcal{M}\times\mathcal{J}} p_{ij}))$. $\square$

2.3 A Better Bound for the Restricted Assignment Problem

In this section we consider the restricted version of our problem, where $p_{ij} \in \{p_j, \infty\}$, for each job $j \in \mathcal{J}$, and each machine $i \in \mathcal{M}$. This subclass is NP-hard to approximate within a factor better than $\frac{3}{2}$, which is also the best known lower bound for the general version [LST90]. In the restricted version, any job j with processing time $p_{ij} < \infty$ is feasible on machine i, since $p_j \leq T_{opt}$ for every $j \in \mathcal{J}$. Hence, the feasibility parameter of a restricted instance $\mathcal{I}$ is exactly $\varphi(\mathcal{I}) = \frac{min_{j \in \mathcal{J}} |\{i \in \mathcal{M} : p_{ij} < \infty\}|}{m}$, which can be computed efficiently. Fully-feasible restricted instances correspond to the identical machines instances, hence we do not consider especially fully-feasible instances in this case. Also, we say that an assignment $\sigma : \mathcal{M} \to \mathcal{J}$ is *feasible* if $p_{ij} < \infty$ for every machine i and job j such that $\sigma(j) = i$.

For this variant, we show that a better bound than in Theorem 2.4 can be achieved by a much simpler and more efficient combinatorial algorithm, and for every feasibility parameter. Denote by $p_{max} = max_{j \in \mathcal{J}} p_j$ the largest processing time of some restricted instance $\mathcal{I}$. Gairing et al. [GL+04] presented a $(2 - \frac{1}{p_{max}})$-approximation algorithm for the restricted assignment problem. Using techniques from [GL+04], we obtain an approximation algorithm which yields an assignment of makespan at most $p_{max} + \frac{L_{opt}}{\varphi}$, where φ, is the feasibility parameter of the instance.

Overview of the Algorithm of Gairing et al. We describe below the procedure $\mathcal{UBF}$, used in [GL+04]. Let $\mathcal{I}$ be an instance for the restricted assignment problem. Let Δ be an integer that will be determined by binary search, to be a lower bound on T_{opt}.

Let σ be a feasible assignment and let $G_\sigma = (W \cup V, E_\sigma)$ be a directed bipartite graph where $W = \{w_j : j \in \mathcal{J}\}$ consists of the job nodes, and $V = \{v_i : i \in \mathcal{M}\}$ consists of the machine nodes. For any job node w_j and any machine node v_i, if $\sigma(j) = i$ there is an arc in E_σ oriented from v_i to w_j; if $\sigma(j) \neq i$ and j is feasible on machine i, then there is an arc in E_σ oriented from w_j to v_i.

Given a feasible assignment σ, consider the partition of machines into three subsets: $\mathcal{M}^+(\sigma)$ (overloaded), $\mathcal{M}^-(\sigma)$ (underloaded), and $\mathcal{M}^0(\sigma)$ (all the remaining machines). A machine $i \in \mathcal{M}^+(\sigma)$ is overloaded if $load_\sigma(i) \geq p_{max} + \Delta + 1$. A machine $i \in \mathcal{M}^-(\sigma)$ is underloaded if $load_\sigma(i) \leq \Delta$. The remaining machines, which are neither overloaded nor underloaded, form the set $\mathcal{M}^0(\sigma) = \mathcal{M} \setminus (\mathcal{M}^-(\sigma) \bigcup \mathcal{M}^+(\sigma))$.

The procedure $\mathcal{UBF}(\sigma, \Delta)$ starts with some initial feasible assignment of jobs to machines and iteratively improves the makespan until it obtains an assignment with makespan of $p_{max} + \Delta$, or declares that an assignment of makespan Δ does not exist. In each iteration, the algorithm finds an augmenting path in G_σ, from an overloaded machine to an underloaded machine, and pushes jobs along this path, by performing a series of job reassignments between machines on that path. This results in balancing the load over the machines, i.e., reducing the load of the source that is an overloaded machine, and increasing the load of the destination that is an underloaded machine, while preserving the load of all other machines. Figure 2.2 gives a pictorial example of this operation.

$\mathcal{UBF}$ terminates when there is no path from an overloaded machine to an underloaded machine in G_σ, and this occurs after $O(mS)$ steps, where $S = \sum_{j \in \mathcal{J}} |\{i : p_{ij} < \infty\}|$. Let τ be the resulting assignment after $\mathcal{UBF}(\sigma, \Delta)$ terminates. Then, it is shown in [GL+04] that if $\mathcal{M}^+(\tau) \neq \emptyset$, then $T_{opt} > \Delta$.

The procedure $\mathcal{UBF}$ combined with a binary search over the possible range of values for Δ, is used to identify the smallest Δ such that a call to $\mathcal{UBF}(\sigma, \Delta)$ returns an assignment τ with $\mathcal{M}^+(\tau) = \emptyset$. This yields the approximation ratio of $2 - \frac{1}{p_{max}}$.

The running time of the approximation algorithm is factored by a value that is logarithmic in the size of the range in which we search for Δ, e.g., $[0, \sum_{j \in \mathcal{J}} p_j]$. Thus, the algorithm of [GL+04] computes an assignment having makespan within a factor of $2 - \frac{1}{p_{max}}$ from the optimal in time $O(mSlogP)$, where $P = \sum_{j \in \mathcal{J}} p_j$.

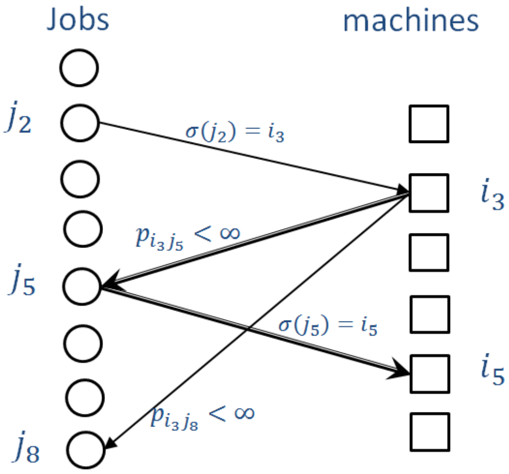

Figure 2.2: The bipartite graph G_σ. By changing the orientation of the path $i_3 \to j_5 \to i_5$, we remove job j_5 from i_5 and schedule it on machine i_3.

2.3.1 Approximation Algorithm

Let $\mathcal{I}$ be an instance of the restricted assignment problem. The feasibility parameter of $\mathcal{I}$ is exactly $\varphi = \varphi(\mathcal{I}) = \frac{min_j |\{i : p_{ij} < \infty\}|}{m}$. Let σ be an initial feasible assignment for $\mathcal{I}$, and consider the bipartite graph G_σ.

Note that any feasible assignment σ for an instance of the restricted assignment problem has an average machine load

$$L(\sigma) = \frac{1}{m} \sum_{j \in \mathcal{J}} p_j \tag{2.1}$$

Thus, $L_{opt} = \frac{1}{m} \sum_{j \in \mathcal{J}} p_j$ for any instance $\mathcal{I}$.

Our algorithm proceeds as follows.

Algorithm 2 $\mathcal{A}_{RES}(\mathcal{I})$

```
1. Fix $\Delta = \lfloor \frac{L_{opt}}{\varphi(\mathcal{I})} \rfloor$.

2. Apply $\mathcal{UBF}(\sigma, \Delta)$ and return the resulting assignment.
```

Theorem 2.10. *For any instance $\mathcal{I}$ with feasibility parameter $\varphi(\mathcal{I})$, Algorithm 2 yields a schedule of makespan at most $p_{max} + \frac{L_{opt}}{\varphi(\mathcal{I})} = p_{max} + \frac{\sum_{j \in \mathcal{J}} p_j}{m \cdot \varphi(\mathcal{I})}$, in time $O(m^2 n)$.*

Corollary 2.11. *Let $\mathcal{I}$, be an instances with feasibility parameter $\varphi(\mathcal{I}) = \frac{d}{m}$, for some $d = 1, ..., m$. Then Algorithm 2 yields a schedule of makespan at most $p_{max} + \frac{\sum_{j \in \mathcal{J}} p_j}{d}$. Therefore, for instances with sufficiently large d (which is equivalent to a large feasibility parameter), namely d such that $d \cdot p_{max} > \sum_{j \in \mathcal{J}} p_j$, we get a schedule with a makespan better than twice the optimal makespan. Let $\frac{3}{2} < r < 2$, then for instances with feasibility parameter $\frac{d}{m}$ such that $d \cdot (r-1) p_{max} > \sum_{j \in \mathcal{J}} p_j$, Algorithm 2 is a r-approximation algorithm.*

The correctness of Theorem 2.10 follows from the next lemma.

Lemma 2.12. *Let $\mathcal{I}$ be an instance with feasibility parameter $\varphi(\mathcal{I})$. Let σ be an initial feasible assignment for $\mathcal{I}$. Then $\mathcal{UBF}(\sigma, \lfloor \frac{L_{opt}}{\varphi(\mathcal{I})} \rfloor)$ terminates with $\mathcal{M}^+ = \emptyset$.*

Proof. Let τ be the assignment when $\mathcal{UBF}(\sigma, \lfloor \frac{L_{opt}}{\varphi(\mathcal{I})} \rfloor)$ terminates. At this point,there is no path from a machine in $\mathcal{M}^+(\tau)$ to a machine in $\mathcal{M}^-(\tau)$ in the graph G_τ. Assume that $\mathcal{M}^+(\tau) \neq \emptyset$. Then there exists a machine i' with $load_\tau(i') \geq p_{max} + \lfloor \frac{L_{opt}}{\varphi} \rfloor$. Denote by $\mathcal{M}_{i'}$ the set of machines i such that v_i is reachable from $v_{i'}$, then $\mathcal{M}_{i'} = \{i \in M :$ there is a directed path in G_τ from $v_{i'}$ to v_i $\}$.

Obviously, there exists a job j', such that $\tau(j') = i'$. Thus, there is an edge $(v_{i'}, u_{j'})$ in G_τ. Since φ is the feasibility parameter of $\mathcal{I}$, there exists at least φm machines on which $u_{j'}$ is feasible, i.e., there exists an edge from $u_{j'}$ to each one of these machines. By appending each of these edges to $(v_{i'}, u_{j'})$ we get a directed path from $v_{i'}$ to at least φm machines (including $v_{i'}$). Therefore, we conclude that $|\mathcal{M}_{i'}| \geq \varphi m$.

Now, we compute a lower bound on the average machine load for τ, by summing the loads of all the machines $i \in \mathcal{M}$. We have that $i' \in \mathcal{M}^+(\tau)$, thus $load_\tau(i') > p_{max} + \frac{L_{opt}}{\varphi}$. Also, there is no path from $v_{i'}$ to machines in $\mathcal{M}^-(\tau)$, and therefore $\mathcal{M}_{i'} \cap \mathcal{M}^-(\tau) = \emptyset$. Thus, for all $i \in \mathcal{M}_{i'}$ it holds that $load_\tau(i) \geq \lfloor \frac{L_{opt}}{\varphi} \rfloor + 1$. Hence,

$$\begin{aligned}
\sum_{i \in \mathcal{M}} load_\tau(i) &\geq load_\tau(i') + \sum_{i \in \mathcal{M}_{i'}, i \neq i'} load_\tau(i) \\
&> p_{max} + \lfloor \tfrac{L_{opt}}{\varphi} \rfloor + 1 + (|\mathcal{M}_{i'}| - 1) \cdot (\lfloor \tfrac{L_{opt}}{\varphi} \rfloor + 1) \\
&= p_{max} + |\mathcal{M}_{i'}|(\lfloor \tfrac{L_{opt}}{\varphi} \rfloor + 1) \\
&\geq p_{max} + |\mathcal{M}_{i'}|((\tfrac{L_{opt}}{\varphi} - 1) + 1) \\
&\geq p_{max} + m\varphi(\tfrac{L_{opt}}{\varphi}) \\
&= p_{max} + mL_{opt} \\
&> mL_{opt}
\end{aligned}$$

We have shown that the sum of loads of the assignment τ is greater than mL_{opt}. Hence, the average load for τ is greater than L_{opt}. A contradiction. By 2.1, the average load of any schedule, and in particular τ, cannot exceed L_{opt}. $\square$

Proof of Theorem 2.10. Let $\mathcal{I}$ be an instance with feasibility parameter φ. Let σ be some initial feasible assignment. By Lemma 2.12, when $\mathcal{UBF}(\sigma, \lfloor \frac{L_{opt}}{\varphi} \rfloor)$ terminates, $\mathcal{M}^+ = \emptyset$. Therefore, the maximum load of the resulting assignment is at most $p_{max} + \lfloor \frac{L_{opt}}{\varphi} \rfloor = p_{max} + \lfloor \frac{\sum_{j \in \mathcal{J}} p_j}{\varphi} \rfloor$. The running time of the algorithm equals to the running time of one call to $\mathcal{UBF}$, which is $O(mS)$, where $S = \sum_{j \in \mathcal{J}} |\{i : p_{ij} < \infty\}|$. Since $S \leq m \cdot n$ for every instance $\mathcal{I}$, the algorithm terminates after $O(m^2 n)$ steps. $\square$

Note that our algorithm has better running time than the algorithm of [GL+04], since we use a single call to the procedure $\mathcal{UBF}$, in contrast to the $(2 - \frac{1}{p_{max}})$-approximation algorithm of [GL+04], which uses binary search to find the best value for Δ, resulting in an overall running time of $O(mSlogP)$, where $P = \sum_{j \in \mathcal{J}} p_j$ is the sum of processing times of all jobs.

Chapter 3

Fixed-Parameter Algorithms for Scheduling on Unrelated Machines

3.1 Preliminaries

Some problems can be solved exactly, or approximately, by algorithms that are exponential only in the size of a fixed parameter while polynomial in the size of the input. Such an (approximation) algorithm is called a *fixed-parameter tractable (FPT)* (approximation) algorithm, because the problem can be (approximately) solved efficiently for small values of the fixed parameter.

Problems in which some parameter k is fixed are called parameterized problems. A parameterized problem that allows for such an FPT algorithm is said to be a *fixed-parameter tractable* problem and belongs to the class FPT.

We give some definitions formalizing this concept.

Definition 3.1. *A problem is said to be FPT if it can be solved by an algorithm $\mathcal{A}$ that runs in time $f(k) \cdot poly(|\mathcal{I}|)$, for every instance $\mathcal{I}$ with parameter k, and where f is a function independent of $|\mathcal{I}|$. The algorithm $\mathcal{A}$ is called an FPT algorithm.*

We can similarly define FPT-approximation algorithms.

Definition 3.2. *A problem is said to have an r-FPT approximation algorithm, if there exists an r-approximation algorithm $\mathcal{A}$ to the problem, that runs in time $f(k) \cdot poly(|\mathcal{I}|)$, for every instance $\mathcal{I}$ with parameter k, and where f is a function independent of $|\mathcal{I}|$.*

Definition 3.3. *A family $\{\mathcal{A}_\epsilon : \epsilon > 0\}$, where $\mathcal{A}_\epsilon$ is a $(1+\epsilon)$-FPT approximation algorithm for all $\epsilon > 0$, is called a parametrized approximation scheme.*

3.2 Parametrized Approximation Scheme for Scheduling on Unrelated Machines

Consider the problem of minimizing the makespan on unrelated machines, i.e., scheduling a set $\mathcal{J}$ of n jobs, $j = 1, .., n$, on a set $\mathcal{M}$ of m unrelated parallel machines, $i \in \mathcal{M}$, where each job j has a processing time of p_{ij} on machine i and the objective is to find a schedule with minimum makespan.

Our parameter k of a scheduling instance is the number of machine-job pairs $(i, j) \in \mathcal{M} \times \mathcal{J}$ such that $p_{ij} > \epsilon \cdot T$ for some $\epsilon \in (0, 1]$ and a feasible value T. We will show that by rounding a solution to the MILP formulation of the problem where the k variables x_{ij} such that $p_{ij} > \epsilon \cdot T$ are integral, we can get an assignment with makespan at most $(1 + \epsilon)$ the optimal makespan.

Given a positive integer T, let x_{ij} be an indicator to the assignment of job j on machine i. Consider the following mixed integer linear program.

$$
\begin{aligned}
MILP(\epsilon, T): \qquad & \sum_{i=1}^{m} x_{ij} = 1, && \text{for } j = 1, ..., n \\
& \sum_{j=1}^{n} p_{ij} x_{ij} \le T, && \text{for } i = 1, ..., m \\
& x_{ij} \ge 0, && \text{for } i, j \text{ such that } p_{ij} \le \epsilon \cdot T \\
& x_{ij} \in \{0, 1\}, && \text{for } i, j \text{ such that } p_{ij} > \epsilon \cdot T \\
& x_{ij} = 0, \ if \ p_{ij} > T, && \text{for } i = 1, ...m \, , j = 1, ..., n
\end{aligned}
$$

One can see that integer solutions to the above MILP are in one to one correspondence with assignments $\sigma : \mathcal{J} \to \mathcal{M}$ of makespan at most T, and that any feasible solution to $MILP(\epsilon, T)$ has the property that the variables x_{ij}, such that $p_{ij} > \epsilon T$, are integral.

Theorem 3.1. *Let $\mathcal{I}$ be a scheduling instance, let $\epsilon \in [0, 1]$ and let T be a positive integer such that $MILP(T, \epsilon)$ is feasible. Then an assignment of makespan at most $(1 + \epsilon)T$ can be found in time $2^k \cdot poly(|\mathcal{I}|)$, where $k = |\{(i, j) : p_{ij} > \epsilon \cdot T\}|$.*

Proof. Let x_{ij}, $i = 1, ..., m$, $j = 1, ..., n$ be a solution to $MILP(\epsilon, T)$.

Then for every pair $(i, j) \in \mathcal{M} \times \mathcal{J}$ such that $p_{ij} > \epsilon \cdot T$, the corresponding x_{ij} is either 0 or 1. Let l_i be the number of jobs j for which $p_{ij} > \epsilon \cdot T$ on machine i and such that the corresponding variable x_{ij} equals to 1. Recall the rounding technique as in Theorem 2.1. Then the first l_i slots of machine i are full and therefore the capacity left in the slots for the other jobs, with $p_{ij} \le \epsilon \cdot T$ is $T - \sum_{j : p_{ij} > \epsilon T} p_{ij}$. From Theorem 2.1, we get that the jobs that are assigned by the rounding to the remaining slots $l_i + 1, ..., k_i$ (where $k_i = \left\lceil \sum_{j \in \mathcal{J}} x_{ij} \right\rceil$) contribute at most $p_{max}^{l_i+1} + T - \sum_{j : p_{ij} > \epsilon T} p_{ij}$ to the load of machine i. Therefore the load of machine i is at most $\sum_{j:p_{ij}>\epsilon T} p_{ij} + \left(p_{max}^{l_i+1} + T - \sum_{j:p_{ij}>\epsilon T} p_{ij}\right) = p_{max}^{l_i+1} + T$. Now, $p_{max}^{l_i+1} \leq \epsilon \cdot T$, and therefore the total load of machine i is at most $(1+\epsilon)T$.

The algorithm runs in time $2^k \cdot poly(|\mathcal{I}|)$, since the bottle-neck of the algorithm is obtaining a feasible solution to the MILP. This is done by brute-force search of at most 2^k possible binary values for all the variables x_{ij} corresponding to $(i,j) \in S_\epsilon$, by fixing them and finding a solution for the resulting LP which can be done in polynomial-time in $|\mathcal{I}|$ [K84].

□

3.3 An FPT Algorithm for Graph-Balancing

In this section we consider the special case of the restricted assignment problem, where each job can be assigned to at most two machines, with the same processing time on either machine. For this special case, Ebenlendr et al.[EKS08] presented a 1.75-approximation algorithm for the minimum makespan problem.

An instance of the scheduling problem can be modeled as an undirected, multi-graph, with m nodes (a node for each machine) and n edges, such that every job j is associated with an edge of weight p_j between both machine nodes on which it can be processed, or a loop of weight p_j on the only machine node on which it can be processed. Minimizing the makespan is then equivalent to the problem of *graph-balancing*, i.e., of orienting each edge, such that the maximum weighted in-degree over all nodes is minimized. We exploit this graph representation of the problem to develop an FPT algorithm for this case.

Definition 3.4. *The maximum degree r of an undirected graph $G = (V, E)$ is the maximum number of neighbors of any vertex, i.e., $r = max_{v \in V}|N(v)|$, where $N(v) = \{e \in E : vu, for\ some\ u \in V\}$.*

We give below an FPT algorithm for graph balancing, where the parameters are the width of the tree decomposition of the graph, and the maximum degree of the graph.

Note that we cannot hope for obtaining an FPT algorithm with the fixed parameter being only the maximum degree of the graph, as from the hardness proof for general instance [EKS08], if follows that he problem is hard to approximate within a factor less than $\frac{3}{2}$ even on bounded degree graphs, i.e., when the maximum degree is some constant.

Intuitively, a tree decomposition represents the vertices of a given graph G as subtrees of a tree, in such a way that vertices in the given graph are adjacent only when the corresponding subtrees intersect.

Definition 3.5. *Given a graph $G = (V, E)$, a tree decomposition is a pair $\langle X, T\rangle$, where $X = \{X_1, ..., X_t\}$ is a family of subsets of V (also called bags), and T is a tree whose nodes are the subsets X_i, satisfying the following properties:*

1. *The union of all sets X_i equals V. That is, each graph vertex is associated with at least one tree node.*

2. *For every edge (v, w) in the graph, there is a subset X_i that contains both v and w. That is, vertices are adjacent in the graph only when the corresponding subtrees have a node in common.*

3. *If X_i and X_j both contain a vertex v, then all nodes X_k of the tree in the (unique) path between X_i and X_j contain v as well. That is, the nodes associated with vertex v form a connected subset of T. It can be stated equivalently that if X_i, X_j and X_k are nodes, and X_k is on the path from X_i to X_j, then $X_i \cap X_j \subseteq X_k$.*

The width of a tree decomposition is the size of its largest set X_i minus one. The treewidth $tw(G)$ of a graph G is the minimum width among all possible tree decompositions of G.

It is observed in [B88, N06], that many algorithmic problems that are NP-complete for arbitrary graphs can be solved efficiently by dynamic programming for graphs of bounded treewidth, using the tree decompositions of these graphs. We show that the problem of graph balancing can be solved efficiently by dynamic programming for graphs of bounded treewidth and bounded vertex degree (the maximum number of neighbors of a vertex).

Theorem 3.2. *Let G be an undirected edge weighted multi-graph, with given tree decomposition $\langle\{X = \{X_1, ..., X_t\}, T\rangle$ of width w. Then the graph balancing problem parameterized by the graph treewidth and the maximum degree, r, is solvable in time $O(2^{2w\cdot r} \cdot wr \cdot |X|$.*

Proof. We show how the problem can be solved using dynamic programing on the tree decomposition of G. The idea is to examine for each bag $X_i \in X$ all the possibilities of feasible assignments to the machines represented by the vertices in the bag X_i, and the jobs represented by the edges of the subgraph $G[X_i]$, induced by the vertices in bag X_i. This information is stored in a table B_i corresponding to each bag X_i. The tables will be updated in a post-order manner, starting at the leaves of the tree decomposition and ending at the root. During this update process, it is guaranteed that local solutions for each subgraph corresponding to a bag of the tree decomposition are combined into a global optimal solution for the overall graph G.

The algorithmic details are as follows.

Step 0: Set an initial orientation on the edges of the graph G. For each bag $X_i = \{v_1^i, ..., v_{n_i}^i\}$, $|X_i| = n_i$, let $E[X_i] = \{e_1^i, ..., e_{m_i}^i\}$, $|E_{G[X_i]}| = m_i$. Compute the following table of 2^{m_i} rows, and $m_i + n_i + 1$ columns:

$$B_i = \begin{array}{cccc||ccccc}
e_1^i & e_2^i & \cdots & e_{m_i}^i & l_i(v_1^i) & l_i(v_2^i) & \cdots & l_i(v_{n_i}^i) & T_i() \\
\hline
0 & 0 & \cdots & 0 & & & & & \\
0 & 0 & \cdots & 1 & & & & & \\
0 & 0 & \cdots & 1 & & & & & \\
\vdots & \vdots & \vdots & \vdots & & & & & \\
1 & 1 & \cdots & 0 & & & & & \\
1 & 1 & \cdots & 1 & & & & &
\end{array}$$

The table consists of $2^{n_i r}$ rows and $n_i r + 1$ columns. Each row represents an assignment to the sub-problem induced by the subgraph $G[X_i]$. Each row is a 0-1 sequence of length $n_i r$ that determines which of the edges in $G[X_i]$ is directed oppositely than its direction in the given initial orientation (1 if it is the opposite orientation, and 0 otherwise). Formally, we can describe an assignment by a mapping

$$A_i : E[X_i] = \{e_{i_1}, ..., e_{i_{m_i}}\} \to \{0, 1\}.$$

Given the mapping A_i, let $In(v_j)$ denote the set of incoming edges for v_j, for $v_j \in X_i$. The last column, $T_i()$, is the makespan of the assignment A_i.

Step 1: Table initialization.

For every table B_i and assignment $A_i : E[X_i] \to \{0, 1\}$ set

$$l_i(v_j^i)(j) = \sum_{e \in In(v_j)} c(e)$$

for every $v_{i_1}, ..., v_{i_{n_i}} \in X_i$, where $c(e)$ is the cost of edge e, which corresponds to the processing time of the job associated with the edge e.

Step 2: Dynamic programming.

We now go through the tree decomposition of G, from the leaves to the root, and compare the corresponding tables against each other. Let i be the parent node of j. We show how the table for X_i can be updated by the table for X_j. Assume that $X_i = \{u_1, .., u_s, v_1^i, ..., v_{n_i - s}^i\}$ and $X_j = \{u_1, .., u_s, v_1^j, ..., v_{n_j - s}^j\}$, and that $E[X_i] = \{e_1, .., e_t, e_1^i, ..., e_{m_i - t}^i\}$ and $E[X_j] = \{e_1, .., e_t, e_1^j, ..., e_{m_j - t}^j\}$.

For each assignment $A : \{e_1, ..., e_t\} \to \{0, 1\}$, and each extension $A_i : E[X_i] \to \{0, 1\}$ of A, we consider an assignment A_j, which is an extension of A, that minimizes the new makespan, i.e., for each assignment A_j, which is an extension of A, we calculate

$$l_i^j(u_k) = l_j(u_k) + l_i(u_k) - \sum_{e \in E[X_i] \cap E[X_j] \wedge e \in In(u_k)} c(e),$$

for $k = 1, ..., s$. Then, we calculate the makespan

$$T_i^j = max\left[T_j(A_j), max_{1 \le k \le s} l_i^j(u_k), max_{1 \le k \le n_i - s} l_i(v_k^i)\right]$$

We update the entry for A_i with l_i^j and $T_i^{\hat{j}}$ for $\hat{j} = argmin(T_i^j)$.

The values of $l_i(v)$, $v \in X_i$, and $T_i()$ grows by the minimal value for the makespan of the assignment problem induced by all the vertices contained in the subtree rooted at node i. If i has several children $j_i, ..., j_l$, then table B_i is updated successively against all tables $B_{j_1}, ..., B_{j_l}$ in the same way. All this is repeated until the root node is finally updated.

Step 3: Construction of a minimum makespan assignment.

The length of a minimum makespan assignment is derived from the minimal entry of the last column, $T_i()$, of the root node table. The assignment of the corresponding row shows where to assign the jobs represented by the edges in the subgraph induced by the vertices of the root bag. By recording in Step 2 how the respective minimum of each bag was determined by its children, one can easily determine the assignment of all edges in the graph.

This concludes the description of the dynamic programming algorithm. It remains to show its correctness and running time.

1. The first and second conditions in Definition 3.5, namely $V = \cup_X X_i$ and $\forall e \in E \; \exists X_i \in X : e \in E[X_i]$, guarantee that every machine and job in the instance is considered through the computation.

2. The third condition in Definition 3.5 guarantees the consistency of the dynamic programming. If a vertex $v \in V$ occurs in two different bags X_{i_1} and X_{i_2}, then it is guaranteed that for the computed minimum makespan assignment only one set of jobs can be scheduled on the machine associated with that vertex v.

As for the running time of the algorithm, for each edge (X_i, X_j) in the tree decomposition of G, and for each of 2^{m_i} assignments A_i, we go over all the assignments A_j that agree with A_i on the edges in $E[X_i] \cap E[X_j]$ (at most 2^{m_j}), and do a computation of time $O(n_i + m_i)$. This results in complexity of $O(2^{m_i+m_j}) \cdot (n_i + m_i) \cdot |X|$. Since $m_k \leq n_k \cdot r$ and $n_k \leq w$ for all $X_k \in X$, we have that the complexity of the algorithm is $O(2^{2wr} \cdot wr \cdot |X|)$. $\square$

Chapter 4

Reoptimization Algorithms for Scheduling Problems

4.1 Preliminaries

Let Π_{ID} and Π_{UN} denote the makespan minimization problems on identical and uniform machines, respectively. In the *reoptimization* model developed in [STT12], we consider instances of the scheduling problem that can change dynamically over time. Our goal is to compute assignments within some guaranteed approximation for the new problem instances, derived from the previous instances. Since the transition from one assignment to another incurs some cost (for example, the cost of pausing the execution of a process on one machine and resuming its execution on another), an additional goal is to have the solution for the new instance *close* to the original one (under a certain distance measure).

Let $\mathcal{I}_0 = (\mathcal{M}_0, \mathcal{J}_0)$ be an instance of jobs and machines. Let $m_0 = |\mathcal{M}_0|$ and $n_0 = |\mathcal{J}_0|$ and let $\sigma_0 : \mathcal{J}_0 \to \mathcal{M}_0$ be some initial assignment for $\mathcal{I}_0$. We denote by $\mathcal{I} = (\mathcal{M}, \mathcal{J})$ a new instance derived from $\mathcal{I}_0$ by an admissible operation, e.g., addition or removal of jobs and/or machines. For any job $j \in \mathcal{J}$ and a feasible assignment $\sigma : \mathcal{J} \to \mathcal{M}$, we are given the transition cost of j when moving from the initial assignment σ_0 to σ. We denote this transition cost by $c_{\sigma_0}(j, \sigma)$. The goal is to find an optimal assignment for $\mathcal{I}$, for which the total transition cost, given by $\sum_{j \in \mathcal{J}} c_{\sigma_0}(j, \sigma)$, is minimized.

Recall that, given an optimization problem Π, we denote by $R(\Pi)$ the reoptimization version of Π.

Definition 4.1. *An algorithm $\mathcal{A}$ yields an (r, ρ)-reapproximation for $R(\Pi_{ID})$ ($R(\Pi_{UN})$), for $r, \rho \geq 1$, if for any instance $\mathcal{I}$ for Π_{ID} (Π_{UN}), $\mathcal{A}$ outputs an assignment of makespan at most ρ times the minimal makespan for $\mathcal{I}$, and of total transition cost at most r times the minimal transition cost to an optimal assignment for $\mathcal{I}$.*

We consider below the case where transition costs can take values in $\{0, 1\}$. In particular, job $j \in \mathcal{J}$ incurs a unit transition cost either if (i) $j \in \mathcal{J}_0$ and is moved to a different machine in the schedule for I, or (ii) $j \in \mathcal{J} \setminus \mathcal{J}_0$, i.e., j is assigned to a machine for the first time in the schedule for $\mathcal{I}$. Otherwise, the transition cost for job j is equal to 0. Formally, $c_{\sigma_0}(j, \sigma) = 1$ if $j \notin \mathcal{J}_0$, or if $j \in \mathcal{J}_0$ and $\sigma_0(j) \neq \sigma(j)$; otherwise, $c_{\sigma_0}(j, \sigma) = 0$.

Definition 4.2. *A polynomial time reapproximation scheme (PTRS) for $R(\Pi)$ is an algorithm that, given the inputs $\mathcal{I}_0$ and $\mathcal{I}$ for $R(\Pi)$ and parameters $\epsilon_1, \epsilon_2 \geq 0$, yields a $(1+\epsilon_1, 1+\epsilon_2)$-reapproximation for $R(\Pi)$, in time polynomial in $|\mathcal{I}_0|$ and $|\mathcal{I}|$.*

4.2 A $(1, 1+\epsilon)$-Reapproximation Algorithm for Makespan Minimization on Identical Machines

We present below a reapproximation algorithm, $\mathcal{A}_{ID}$, for the problem of minimizing the makespan on identical machines. The algorithm uses a *relaxed* packing of items in bins, where the items correspond to jobs, and the bins represent the set of machines.

Definition 4.3. *Given a set of bins, each of capacity $K > 0$, and a set of items packed in the bins, we say that the packing is ϵ-* relaxed*, for some $\epsilon > 0$, if the total size of items assigned to each bin is at most $(1+\epsilon)K$.*

4.2.1 Algorithm

Our algorithm for reoptimizaing makespan minimization on identical machines accepts as input the instances $\mathcal{I}_0$ and $\mathcal{I}$, the initial assignment of jobs to the machines, σ_0, and an error parameter $\epsilon > 0$. The algorithm proceeds as follows.

We apply a $(1+\epsilon)$-approximation algorithm [H96, AA+98] on the new instance to obtain a solution of makespan $T \leq (1+\epsilon)C^*_{max}$. Then, we split our instance into large and small jobs, round down the large jobs processing times (to have a polynomial-size collection of feasible configurations of large jobs on the machines), such that the load of each machine does not exceed T. Then, we iterate on this collection in order to find the configuration that minimizes the transition cost from the original solution. We prove that after we inflate the rounded jobs to their original processing times, and greedily assign all the small jobs within the configuration, the resulting makespan is at most $(1+\epsilon)C^*_{max}$.

We give below a detailed description of our algorithm, $\mathcal{A}_{ID}$. Let $C^*_{max}(\mathcal{I})$ denote the minimum makespan for an instance $\mathcal{I}$. For simplicity of the presentation, for the case where $m < m_0$, we assume w.l.o.g. that the omitted machines are $m+1, m+2, \ldots, m_0$.

Algorithm 3 $\mathcal{A}_{ID}(\mathcal{I}_0, \mathcal{I}, \sigma_0)$

```
1. Let $\epsilon_0 = \frac{\epsilon}{4}$. Use a PTAS for makespan minimization on identical machines to find
   $T \leq (1+\epsilon_0)C^*_{max}(\mathcal{I})$.

2. Define $\alpha_j = \frac{p_j}{T}$ for all $j \in \mathcal{J}$, and represent each machine as a bin of unit capacity.
   Consider the jobs as items whose sizes are $\alpha_j \in (0,1]$.

3. An item $j \in \mathcal{J}$ is small if it has a size at most $\epsilon_0$; otherwise, item $j$ is large.

4. Round down the sizes of the large items to the nearest multiple of $\epsilon_0^2$. Denote the
   rounded sizes $\bar{\alpha_j}$, for every large item $j$.

5. For any feasible assignment of rounded large items in the $m$ bins, given by the con-
   figuration $\mathcal{C} = \{C^1, \ldots, C^m\}$, do:

    (i) Let $\ell = \max\{m_0, m\}$. Construct a complete bipartite graph $G = (U, V, E)$, in
        which $V = \{1, \ldots, \ell\}$, and $U = \{C^1, C^2, \ldots, C^\ell\}$. Each vertex $i \in V$ corresponds
        to the initial configuration of machine i, given by $C_0^i = \{j \in \mathcal{J}_0 : \sigma_0(j) = i\}$, for
        $1 \leq i \leq m_0$; if $m_0 < m$ , set $C_0^i = \emptyset$ for all $m_0 + 1 \leq i \leq \ell$. If $m_0 \geq m$, set $C^i = \emptyset$
        for all $m + 1 \leq i \leq \ell$. Define a cost on the edges $(i, C^k)$, for all $1 \leq i, k \leq \ell$, as
        follows.

        (a) Add the cost of large items that appear in $C^k$ but not in the initial config-
            uration $C_0^i$.
        (b) Add to $C^k$ all the small items that appear in $C_0^i$ but not in $C^k$ and then
            omit the largest small items until the total size of $C^k$ does not exceed 1.
            Add the cost of the omitted items. For an empty configuration $C^k$, the cost
            of the edge $(i, C_k)$, for all $1 \leq i \leq \ell$ is equal to 0.

   (ii) Find a minimum cost perfect matching in the bipartite graph.

  (iii) Add to the solution the omitted small items using First-Fit.

6. Choose the solution of minimum cost, and return the corresponding schedule of the
   jobs on the machines.
```

4.2.2 Analysis

Theorem 4.1. *For any* $\epsilon > 0$*, Algorithm* $\mathcal{A}_{ID}$ *yields in polynomial time a* $(1, 1+\epsilon)$*-reapproximation for* $R(\Pi_{ID})$*,*

We prove the theorem using the next lemmas.

Lemma 4.2. *Let* $\mathcal{I} = (\mathcal{M}, \mathcal{J})$ *be an instance of* Π_{ID}*, for which the minimum makespan is* C^*_{max}*, and let* $T \geq C^*_{max}$*. Let* $\alpha_j = \frac{p_j}{T}$*, for all* $j \in \mathcal{J}$*; then,* $\sum_{j \in \mathcal{J}} \alpha_j \leq m$*.*

Proof. We note that the minimum makespan satisfies $C^*_{max} \geq \frac{\sum_{j \in \mathcal{J}} p_j}{m}$. Since $T \geq C^*_{max}$, we have that $\alpha_j = \frac{p_j}{T} \leq \frac{p_j}{C^*_{max}}$ for all $j \in \mathcal{J}$, therefore,

$$m \geq \sum_{j \in \mathcal{J}} \frac{p_j}{T} = \sum_{j \in \mathcal{J}} \alpha_j.$$

□

Lemma 4.3. *Let* $\widetilde{\mathcal{C}}$ *be a feasible assignment of rounded large items on the* m *machines, given by the configuration* $C = \{C^1, \ldots, C^m\}$*, to which we add in each bin the small items that were not omitted in Step 5 of* $\mathcal{A}_{ID}$*. Then* $\widetilde{\mathcal{C}}$ *can be expanded in polynomial time to an* ϵ_0*-relaxed packing of all items in the input* $\mathcal{I}$*.*

Proof. Let $C = \{C^1, C^2, ..., C^m\}$ be a feasible configuration of the large rounded items, and let $S^1, S^2, ..., S^m$ be the subsets of small items added in Step 5 to the bins, to form $\widetilde{\mathcal{C}}$. Then $\widetilde{\mathcal{C}}$ yields a feasible packing, i.e., for all $1 \leq i \leq m$,

$$\sum_{j \in C^i} \bar{\alpha}_j + \sum_{j \in S^i} \alpha_j \leq 1.$$

Now, since $\bar{\alpha}_j \geq \epsilon_0$ for all $j \in C^i$, the number of large items in bin i is bounded by $\lfloor 1/\epsilon_0 \rfloor$. Also, since $\alpha_j - \bar{\alpha}_j \leq \epsilon_0^2$, for all $j \in L$, we have that

$$\sum_{j \in C^i \cup S^i} \alpha_j \leq \sum_{j \in C^i} (\bar{\alpha}_j + \epsilon_0^2) + \sum_{j \in S^i} \alpha_j \leq \sum_{j \in C^i} \bar{\alpha}_j + \left\lfloor \frac{1}{\epsilon_0} \right\rfloor \epsilon_0^2 + \sum_{j \in S^i} \alpha_j \leq 1 + \epsilon_0.$$

Hence, the packing of $\widetilde{\mathcal{C}}$ is ϵ_0-relaxed. Now, we show that the packing remains ϵ_0-relaxed after we add the small items using First-Fit. Let $\alpha_1, \alpha_2, ...\alpha_t$ denote the sizes of the items packed in $\widetilde{\mathcal{C}}$, and let $\alpha_{t+1}, ..., \alpha_n$ be the sizes of the small unpacked items. By Lemma 4.2, we have

$$\sum_{i=1}^{n} \alpha_j \leq m \tag{4.1}$$

We apply First-Fit in the following *relaxed* manner. Consider the next item in the unpacked list. Starting from bin 1, we seek the first bin in which the item can be added, such that the overall size of the items packed in this bin is at most $1 + \epsilon_0$. Let r_i be the total size of the items packed in bin i after adding the small items. Assume that, after we apply First-Fit, some small items remain unpacked (i.e., none of the bins can accommodate these items). Let $i_\ell, ..., i_n$ denote this subset of items. Then, we have that $(1 + \epsilon_0) - r_i < \alpha_j$ for all $j = \ell, \ldots, n$ and $i = 1, \ldots, m$. It follows, that

$$(1 + \epsilon_0)m - \sum_{i=1}^{m} r_i < m\alpha_j \qquad \forall j = \ell, \ldots, n \tag{4.2}$$

By the definition of r_i, and since we packed all items up to $i_{\ell-1}$,

$$\sum_{i=1}^{m} r_i = \sum_{j=1}^{\ell-1} \alpha_j \tag{4.3}$$

From (4.1), we get that

$$\sum_{j=1}^{\ell-1} \alpha_j \leq m - \sum_{j=\ell}^{n} \alpha_j.$$

Then, from (4.3), we have

$$m - \sum_{j=\ell}^{n} \alpha_j \geq \sum_{i=1}^{m} r_i \tag{4.4}$$

From (4.2) and (4.4), it follows that, for all $\ell \leq j \leq n$,

$$m(1+\epsilon_0) - (m - \sum_{j=\ell}^{n} \alpha_j) \leq m(1+\epsilon_0) - \sum_{i=1}^{m} r_i < m\alpha_j,$$

or,

$$m\epsilon_0 + \sum_{j=\ell}^{n} \alpha_j < m\alpha_j. \tag{4.5}$$

From (4.5), and since $\alpha_j \leq \epsilon_0$ for $j = \ell, ..., n$, we have that $m\epsilon_0 + \sum_{j=\ell}^{n} \alpha_j < m\epsilon_0$; thus, $\sum_{j=\ell}^{n} \alpha_j < 0$. A contradiction, since $\alpha_j \geq 0$ for all j.

Hence, the above relaxed implementation of First-Fit packs all of the remaining small items and yields a relaxed-packing of the original instance. □

Lemma 4.4. *Let OPT be an optimal solution, and let $\mathcal{C} = \{C^1, ..., C^m\}$ be the configuration of large items derived from OPT. Then, the cost of this optimal solution is at least the cost of the solution for $\mathcal{C}$ in Step 5 of $\mathcal{A}_{ID}$.*

Proof. $\mathcal{C}$ is an optimal configuration, therefore $\mathcal{C}$ is also a feasible configuration (since $\sum_{j \in C^i} \frac{p_j}{T} \leq \frac{C^*_{max}}{T} \leq 1$). Let $ALG_{\mathcal{C}}$ be the solution for $\mathcal{C}$ the algorithm generates in Step 5. Assume that $Cost\,(OPT) < Cost\,(ALG_{\mathcal{C}})$. Note that the cost for the large items is the same in both solutions, and therefore the difference between the cost is caused by packing small items. We also note that, in the algorithm, the small items we pay for are those that are omitted from their original bin and transfered to a different one. Since the transition costs are 1, it means that the number of omitted small items in OPT is smaller than in $ALG_{\mathcal{C}}$, or in other words, the number of small items that are packed in their original bin in OPT is greater than their number in $ALG_{\mathcal{C}}$. Therefore, there must exist a bin $1 \leq i \leq m$ for which this holds. Consider the small items packed in bin i in each solution. Denote by $q_1, ..., q_s$ the sizes of small original items of bin i that are packed in bin i in both solutions, by $p_1^A, ..., p_k^A$ and $p_1^O, ..., p_l^O$ two distinct sets of small original items of bin i that are packed in $ALG_{\mathcal{C}}$ and in OPT, respectively, but not in both. Since the algorithm chooses to omit the largest small items first, the fact it chose to omit $p_1^O, ..., p_l^O$ but to keep $p_1^A, ..., p_k^A$ means that $p_1^O, ..., p_l^O$ are not smaller than $p_1^A, ..., p_k^A$. In particular, let $p_1^O = min\left\{p_1^O, ..., p_l^O\right\}$, then we have

$$p_1^O = min\left\{p_1^O, ..., p_l^O\right\} \geq max\left\{p_1^A, ..., p_k^A\right\}. \tag{4.6}$$

Since $C^*_{max} \le T$, we get that

$$\sum_{j \in C^i} p_j + \sum_{j=1}^{s} q_j + \sum_{j=1}^{\ell} p_j^O \le \frac{C^*_{max}}{T} \le 1 \tag{4.7}$$

The algorithm chose, in particular, to omit p_1^O, hence

$$\sum_{j \in C_i} p_j + \sum_{j=1}^{s} q_j + \sum_{j=1}^{k} p_j^A \le 1 + \epsilon_0 \tag{4.8}$$

but also

$$\sum_{j \in C^i} pj + \sum_{j=1}^{s} q_j + \sum_{j=1}^{k} p_j^A + p_1^O > 1 + \epsilon_0. \tag{4.9}$$

By the above discussion, we also have that

$$\begin{aligned} \sum_{j \in C^i} p_j + \sum_{j=1}^{s} q_j + \textstyle\sum_{j=1}^{k} p_j^A + p_1^O \quad & \underset{(4.6)}{\le} \quad \sum_{j \in C^i} p_j + \sum_{j=1}^{s} q_j + kp_1^O + p_1^O \\ & \underset{k<\ell}{\le} \quad \sum_{j \in C^i} p_j + \sum_{j=1}^{s} q_j + \ell p_1^O \\ & \underset{(4.6)}{\le} \quad \sum_{j \in C^i} p_j + \sum_{j=1}^{s} q_j + \textstyle\sum_{j=1}^{\ell} p_j^O \\ & \underset{(4.7)}{\le} \quad 1 \end{aligned}$$

From the last inequality, we have a contradiction to (4.9). Hence, we have the statement of the lemma. □

Lemma 4.5. *[EK72] Let $G = (V, U, E)$ be a bipartite graph with $|V| = |U| = n$ and, and let $c : E \rightarrow \mathbb{R}$ be a cost function on the edges. A minimum cost perfect matching, i.e. a perfect matching $M \subseteq E$ for which $\sum_{e \in M} c(e)$ is minimized, can be found in $O(n^3)$ time.*

Now, we are ready to prove the main theorem.

Proof of Theorem 4.1: Let $\mathcal{C} = C^1, ..., C^m$ be the configuration of large items derived from an optimal solution, OPT. Let $ALG_{\mathcal{C}}$ be the packing obtained for $\mathcal{C}$ in Step 5 of the algorithm. By Lemma 4.4, we have

$$Cost\left(ALG_C\right) \le Cost\left(OPT\right). \tag{4.10}$$

Let ALG be the solution the algorithm outputs.

Since $Cost(ALG) = min\left\{Cost(ALG_C) : C \text{ is a legal configuration }\right\}$, we have that $Cost(ALG) \le Cost(OPT)$.

Now, it is easy to see that if we transform ALG to the schedule of the job on the machines (by returning to the original processing times), we get a solution of makespan at most $(1 + \epsilon_0)T$. Since $T \le (1 + \epsilon_0)C^*_{max}$, the algorithm is a $(1 + \epsilon_0)^2$-approximation algorithm. We note that $(1+\epsilon_0)^2 = (1+\frac{\epsilon}{4})^2 = (1+\frac{\epsilon}{2}+\frac{\epsilon^2}{16}) \leq 1+\epsilon$, thus the algorithm yields a $(1+\epsilon)$-approximation to C^*_{max}.

We conclude that $\mathcal{A}_{ID}$ is a $(1, 1+\epsilon)$-reapproximation algorithm for the reoptimization problem.

Now, we prove that $\mathcal{A}_{ID}$ runs in polynomial time.

- In the first step, we run a PTAS for the new instance $\mathcal{I}$, therefore, this part is polynomial in $\mathcal{I}$ (but exponential in $\frac{1}{\epsilon_0}$).
- Steps 2,3 and 4 are clearly polynomial in $|\mathcal{I}|$.
- For Step 5, we first show that the number of feasible configuration of large rounded items is polynomial in $\mathcal{I}$. Recall that the number of large rounded items in each bin can not exceed $\left\lfloor \frac{1}{\epsilon_0} \right\rfloor$. Since the rounded sizes of large items can be of at most $\left\lceil \frac{1}{\epsilon_0} \right\rceil$ different sizes, the number of feasible configurations is at most $R = \binom{\left\lfloor \frac{1}{\epsilon_0} \right\rfloor + \left\lceil \frac{1}{\epsilon_0^2} \right\rceil}{\left\lceil \frac{1}{\epsilon_0^2} \right\rceil}$. There are at most m bins, implying that the number of feasible configurations is at most $\begin{pmatrix} m+R \\ m \end{pmatrix}$ which is polynomial in m but exponential in $\frac{1}{\epsilon_0}$. Now, for each configuration, we construct the bipartite graph in time polynomial in $|\mathcal{I}|$ and by [EK72] we find a perfect minimum cost matching in time $O(k^3)$ where $k = max\left\{m, m'\right\}$, which is polynomial in $|\mathcal{I}|$.
- Packing of the small items in each iteration using First-Fit is also done in time polynomial in $|\mathcal{I}|$ and Step 6 is clearly polynomial in $|\mathcal{I}|$.

4.3 A $(1, 1+\epsilon)$-Reapproximation Algorithm for Makespan Minimization on Uniform Machines

We present below a reapproximation algorithm, $\mathcal{A}_{UN}$, for the problem of minimizing the makespan on uniform machines. The algorithm uses a *relaxed* packing of items in bins, where the items correspond to jobs, and the bins represent the set of machines.

Definition 4.4. *Given a set of m bins, with positive capacities $K_1, K_2, ..., K_m$, and a set of items packed in the bins, we say that the packing is ϵ-* relaxed*, for some $\epsilon > 0$, if for every $1 \leq i \leq m$, the total size of items assigned to each bin is at most $(1+\epsilon)K_i$.*

4.3.1 Algorithm

Our algorithm for reoptimizing the makespan on uniform machines accepts as input the instances $\mathcal{I}_0$ and $\mathcal{I}$, the initial assignment of jobs to the machines, σ_0, and an error parameter $\epsilon > 0$.

In the previous section, it was convenient to convert the problem into a bin packing problem where all bins have equal size. We will consider the conversion to a bin packing problem also in the case of scheduling on unrelated machines, only here the bins will have variable sizes. In the generalization of the equal-size bin case to the variable-size case, we come across a major obstacle. The size of the subintervals in which the pieces are partitioned depends on the size of the bins in which the pieces are to be packed. Moreover, the definition of large and small pieces depends on the size of the bin in which the pieces are to be packed. Hochbaum and Shmoys [HS88] presented a PTAS for the problem of makespan minimization on uniform machines. They constructed, in polynomial time, a layered directed graph, with two nodes designated "initial" and "success", such that there exists a path from "initial" to "success" in the graph if and only if there is a schedule with makespan at most $1+\epsilon$ times the minimal makespan. From this path one can also define the configuration of "medium" and "large" jobs on each machine, and it is guaranteed that the "small" jobs can be scheduled. We add suitable costs to the edges of the layered graph and show that a feasible solution with optimal cost can be obtained by finding the lightest path from "initial" to "success" in the layered graph. It is also guaranteed that if there is a path from "initial" to "success" in the graph, then there is enough space in the bins to add the remaining pieces (e.g., using First-Fit), in the bins on which they are considered small. This way, as before, we can greedily pack the small pieces in their original bin (as long as we do not exceed its capacity), and pay only for packing the rest.

Overview of the PTAS of [HS88] Assume that s_1 is the highest speed, and represent each machine i as a bin of size s_i. Normalize the bin sizes by s_1. Consider the jobs as items whose sizes are in $(0, 1]$ (by normalizing them by some upper bound on the minimal makespan). Round down piece sizes in $\left(\epsilon^{k+1}, \epsilon^{k}\right]$ to the nearest multiple of ϵ^{k+2}, for some integer $k \geq 0$.

For a bin of size $s_i \in \left(\epsilon^{k+1}, \epsilon^{k}\right]$ define:

- Pieces of sizes in $\left(\epsilon^{k+1}, \epsilon^{k}\right]$ are *large* for the bin.
- Pieces of sizes in $\left(\epsilon^{k+2}, \epsilon^{k+1}\right]$ are *medium* for the bin.
- Pieces of sizes less than or equal to ϵ^{k+2} are *small* for the bin.

For convenience, the interval $\left(\epsilon^{k+1}, \epsilon^{k}\right]$ is referred as interval k. For pieces in interval k, a bin is *large* if it is in interval k, *huge* if it is in interval $k-1$, and *enormous* if it is in intervals $0, ..., k-2$.

A directed layered graph, $G_L = (V_L, E_L)$, is then constructed, in which each node is labeled with a state vector describing the remaining pieces to be packed as large or medium pieces. The graph is grouped into stages, where a stage specifies the large and medium pack of bins in one interval $\left(\epsilon^{k+1}, \epsilon^{k}\right]$. Each layer within a stage corresponds to packing a bin in the corresponding interval. Both the bins within the stage and the stages are arranged in order of decreasing bin size. The state vector associated with each node is of the form $(L; M; V_1, V_2, V)$, where L and M are vectors, each describing a distribution of pieces in the subintervals of $\left(\epsilon^{k+1}, \epsilon^{k}\right]$ and $\left(\epsilon^{k+2}, \epsilon^{k+1}\right]$ respectively. There are two nodes designated "initial" and "success", such that "initial" is connected to the initial state vectors of the first stage, and every final state vector of the final stage is connected to "success". A path from "initial" to "success" in G_L specifies a packing of the rounded medium and large pieces for every bin.

We note that, after the large and medium pack of the bins in interval k, we must allow for the packing of the remaining pieces in interval k that will be packed as *small.* These pieces must be packed in enormous bins for them, therefore, we need to have sufficient unused capacity in the enormous bins to at least contain the total size of these unpacked pieces. This is represented by the value V_1 in the state vector; it records the slack, or unused capacity in the partial packing of the enormous bins with large and medium pieces. For stages corresponding to intervals greater than k, we also need to have the unused capacity in the huge and large bins, and this is the role of V_2 and V, respectively. Since we must represent the possible values in some compact way, we consider the sizes of the pieces that will be packed as small pieces into this as-yet-unused capacity. For V_1, pieces in interval k are small, and thus all pieces to be packed into this unused capacity have rounded sizes that are multiples of ϵ^{k+2}; as a result, it will be sufficient to represent V_l as an integer multiple of ϵ^{k+2}. Similarly, V_2 and V will be represented as integer multiples of ϵ^{k+3} and ϵ^{k+4}, respectively.

The following lemmas, due to [HS88], will be useful in analyzing our algorithm.

Lemma 4.6. *Given $\epsilon > 0$, the layered graph G_L has $O(2m \cdot n^{2/\epsilon^2+3} \cdot 1/\epsilon^6)$ nodes and $O(2m(n/\epsilon^2)^{(2/\epsilon^2)+3})$ edges. The number of nodes in each layer, which is the number of state vectors corresponding to packing the bins of that layer, is $O(n^{2/\epsilon^2}(n/\epsilon^2)^3)$.*

Lemma 4.7. *For any $\epsilon > 0$, there is a one to one correspondence between paths from "initial" to "success" in the layered graph G_L and ϵ-relaxed packings.*

Lemma 4.8. *Given an ϵ-relaxed packing of rounded piece sizes, for some $\epsilon > 0$, restoring the piece sizes to their original size yields a $(2\epsilon + \epsilon^2)$-relaxed packing of the original pieces.*

We give below a detailed description of our algorithm, $\mathcal{A}_{UN}$. Let $C^*_{max}(\mathcal{I})$ denote the minimum makespan for an instance $\mathcal{I}$.

Algorithm 4 $\mathcal{A}_{UN}(\mathcal{I}_0, \mathcal{I}, \sigma_0)$

```
1. Let $\epsilon > 0$. Use a PTAS for makespan minimization on uniform machines to find $T < (1+\epsilon)C^*_{max}(\mathcal{I})$.

2. Let $s_1$ be the largest speed, and assume that $\frac{s_1}{s_m}$ is bounded from above by some given constant $b \geq 1$. Normalize all the processing times by $s_1 \cdot T$, and represent each machine as a bin of capacity $\frac{s_i}{s_1} \leq 1$. Consider the jobs now as pieces with sizes in $(0, 1]$. Denote the new bin sizes by $\frac{1}{b} \leq s_m \leq s_{m-1} \leq ... \leq s_1 = 1$, and denote the piece sizes by $p_1, p_2, ..., p_n$.

3. Round down the piece sizes $p_j \in \left(\epsilon^{k+1}, \epsilon^k\right]$ to the nearest multiple of $\epsilon^{k+2}$.

4. Construct the directed layered graph $G_L$, and define edge costs as follows. Consider the edge $e^k_\ell$ connecting a node of the $(\ell-1)$'th layer to a node of the $\ell$'th layer in some stage that corresponds to interval $k$. $e^k_\ell$ describes the large and medium pack of the $\ell$'th bin of stage $k$. Let ${C_0}^k_\ell$ be the set of pieces packed in that bin in the initial solution. Fix the cost on that edge to be the number of large and medium pieces that appear in ${C_0}^k_\ell$, but are not packed in the bin by $e^k_\ell$. All other edges in $G_L$ are assigned the cost zero.

5. For every choice of exactly one update arc, at the end of each stage, do:

    (i.) Find the lightest path from "initial" to "success" in the graph (if one exists). Define the corresponding partial solution consisting of pieces that are packed as large or medium.

    (ii.) Add the remaining pieces greedily to the enormous bins for them: for each bin $i$, let $S_{0i}$ be the set of remaining pieces that belong to the bin in the initial solution and that are *small* for this bin. Start packing the pieces in $S_{0i}$ in bin $i$, in increasing order of piece sizes, and stop after the total size of packed pieces exceeds for the first time the bin capacity. Pack all the remaining pieces, as *small*, using First-Fit.

6. Return the solution of minimum cost.
```

4.3.2 Analysis

Theorem 4.9. *For any $\epsilon > 0$, Algorithm $\mathcal{A}_{UN}$ yields in polynomial time a $(1, 1+\epsilon)$-reapproximation for $R(\Pi_{UN})$.*

We prove the theorem using the next lemmas.

Lemma 4.10. *Given the partial packing of medium and large pieces (after Step 5(i) in $\mathcal{A}_{UN}$), packing the remaining pieces greedily as small, in Step 5(ii), incurs the minimal cost for packing these pieces.*

Proof. For the partial packing of medium and large pieces, let R be the set of the remaining pieces. We show that packing the pieces in R in enormous bins for them, by the greedy algorithm given in Step 5(ii), results in the minimal cost for packing R. The greedy algorithm first packs every bin with its original pieces that are small. It sorts them in non-decreasing order by piece size and packs them to the original bin by this order, until the bin capacity is exceeded for the first time (or all pieces are packed). This way we guarantee that we pack to every bin the maximal number of small pieces that were originally packed in it. Thus, the number of pieces, packed as small, *not* in their original bin, is minimized. This also implies a minimum packing cost for R. □

Lemma 4.11. *Let OPT be an optimal solution. Then, the cost of OPT is at least the cost of the solution obtained by $\mathcal{A}_{UN}$.*

Proof. Let $\mathcal{C} = \{C^1, ..., C^m\}$ be the configuration of large and medium items derived from OPT. Since $\mathcal{C}$ is optimal, in particular, it is a truly feasible configuration of large and medium pieces that also leaves enough slackness for packing the remaining pieces as *small*, without exceeding the capacity of the bins. Thus, by Lemma 4.7, there is a path from "initial" to "success" in the layered graph, such that $\mathcal{C}$ is the large and medium pack derived from it.

This path also contains exactly one update arc after each stage. Therefore, our algorithm, in particular, considers this choice of update arcs for which it finds a lightest path in the graph (which clearly exists for this choice). by Lemma 4.10, the cost of the solution output by the algorithm is at most the cost of the lightest path plus the minimal cost of packing all the remaining pieces as *small*, which is clearly at most the cost of OPT. □

Lemma 4.12. *For any partial solution of large and medium rounded pieces, derived from a path from "initial" to "success" in G_L, all the remaining rounded pieces can be packed in enormous bins for them, using the greedy algorithm in Step 5(ii), such that the load on each bin of size s_i is at most $s_i(1 + \epsilon + \epsilon^3)$.*

Proof. As shown in [HS88], the small-pack phase, which is done after every stage, is always successfully completed, since there is an update arc if and only if there is sufficient total slack to accommodate all pieces to be packed as small. Using a similar argument, $\mathcal{A}_{UN}$ is able to pack all the remaining small pieces after packing all the large and medium pieces. Consider the rounded pieces packed into a bin of size s_i, which is in interval k. Focus on the small piece j that, when added to the bin, exhausts the usable slack. Then, piece j is of size less than or equal to ϵ^{k+2}, and before piece j was added, the rounded piece sizes did not exceed $s_i + \epsilon^{k+4}$ (recall that the usable slack is rounded up to a multiple of ϵ^{k+4}). Hence, the bin contains pieces of total (rounded) size at most $s_i + \epsilon^{k+4} + \epsilon^{k+2}$. Therefore, if B_i is the set of pieces packed in bin i, then the sum of the rounded piece sizes in B_i is at most $s_i + \epsilon^{k+4} + \epsilon^{k+2}$ which is at most $s_i + \epsilon s_i + \epsilon^3 s_i = s_i(1 + \epsilon + \epsilon^3)$, since $s_i \geq \epsilon^{k+1}$. □

Lemma 4.13. *The number of different choices for update arcs, one after each stage, is at most*
$O((n^{2/\epsilon^2}(n/\epsilon^2)^3)^S)$*, where* S *is the number of stages in the graph.*

Proof. Between every two stages, there are $O(n^{2/\epsilon^2}(n/\epsilon^2)^3)$ update arcs, same as the number of state vectors in each layer (Lemma 4.6). Therefore, the number of possibilities for choosing one arc in each stage is $O((n^{2/\epsilon^2}(n/\epsilon^2)^3)^S)$. $\square$

Lemma 4.14. *The number of stages* S *in* G_L *depends only on* ϵ *and on* b*.*

Proof. The number of stages in G_L is the number of intervals k that contain at least one bin. Since the smallest bin size satisfies $s_m \geq \frac{1}{b}$, S is at most $min\{k \in \mathbb{N} : \epsilon^{k+1} < \frac{1}{b}\}$. $\square$

Now, we are ready to prove the main theorem.

Proof of Theorem 4.9: By Lemma 4.11, taking the best solution among all lightest paths (one for each choice of update arcs), we have a solution whose transition cost is at most the transition cost of an optimal solution. By Lemma 4.12, this solution is also an $(\epsilon+\epsilon^3)$-relaxed packing of the rounded pieces, and by Lemma 4.8, after inflating the pieces to their original sizes, we get a $(2(\epsilon+\epsilon^3)+(\epsilon+\epsilon^3)^2)$-relaxed packing. Hence, by fixing a suitable initial ϵ, e.g., $\epsilon = \epsilon_0/8$, we conclude that $\mathcal{A}_{UN}$ is a $(1, 1+\epsilon_0)$-reapproximation algorithm for our reoptimization problem.

We note that $\mathcal{A}_{UN}$ runs in polynomial time. Constructing the graph $G_L = (V_L, E_L)$ can be done in time linear in the graph size, which is $V_L = O(2m \cdot n^{2/\epsilon^2+3} \cdot 1/\epsilon^6)$ and $E_L = O(2m(n/\epsilon^2)^{(2/\epsilon^2)+3})$, by Lemma 4.6. By Lemma 4.13, the number of different choices of update arcs, is $X = O((n^{2/\epsilon^2}(n/\epsilon^2)^3)^S)$, where S is the number of stages in the graph. By Lemma 4.14, S depends only on ϵ and b. Therefore, the complexity of running e.g., Dijkstra's algorithm [D59] for finding a lightest path from "initial" to "success" and then greedily packing the small pieces, for every choice of update arcs, is $O(X \cdot (Dijkstra(G_L) + Greedy(n,m)))$. Overall, we get that the complexity of the algorithm is $O(X \cdot (Dijkstra(G_L) + Greedy(n,m)))$, which is polynomial in the input size.

Chapter 5

Conclusions and Future Work

5.1 Summary

In this work we studied the fundamental problem of makespan minimization on parallel machines. For the unrelated machines model, we derived an improved bound, which depends on the minimum average machine load and the feasibility parameter of the instance. Our bound is strictly smaller than 2, the best known general upper bound, for a natural subclass of instances.

We further studied the power of parameterization and presented an FPT algorithm and a parameterized approximation scheme for instances that are known to be hard to approximate within factor $\frac{3}{2}$, based on classical complexity theory.

Finally, we initiated the study of the reoptimization variants of makespan minimization on identical and uniform machines, and derived approximation ratios that match the best known ratios in these models.

5.2 Future Work

Our study leaves open several avenues for future research.

Makespan minimization on unrelated machines While reducing the gap between the general lower bound of $\frac{3}{2}$ and the upper bound of 2 remains a prominent open problem, it would be interesting to tighten these bounds for other non-trivial subclasses of instances, either by a constant, or by a function of the input parameters.

Another interesting direction is to study instances with a wider range of feasibility parameters. In this work, we explored instances having large feasibility parameter. On the other hand, Ebenlendr et al. [EKS08] considered instances in which the feasibility parameter is at most $\frac{2}{m}$. Any results on other, intermediate values of this parameter, would shed more light on its role in obtaining better schedules.

Other natural questions are: “Can we obtain better bound for fully-feasible instances, when $T_{opt} = L_{opt}$?” “Can we derive a better lower bound for general instances for which $T_{opt} = L_{opt}$?”

FPT algorithms for scheduling The problem of minimizing the makespan on unrelated machines is NP-hard even if the number of machines or number of jobs is taken as a parameter [LST90], therefore no FPT algorithm exists with only these choices of parameters. In addition, we cannot hope for obtaining an FPT algorithm for graph balancing with the fixed parameter being only the maximum degree of the graph. Indeed, by the hardness proof of [EKS08], the problem is hard to approximate within a factor less than $\frac{3}{2}$ even on bounded degree graphs, i.e., when the maximum degree is some constant. The question whether we can find an FPT algorithm for graph balancing with the treewidth being the only parameter remains open.

Reoptimization in scheduling We list some of the questions arising from our results. "Can we obtain reapproximation algorithm with the same performance guarantees for general instances in the uniform machines model?" "Can we extend our results to arbitrary transition costs?". Finally, "Can we improve the running times of our reapproximation schemes, by adjusting the known EPTASs for identical and uniform machines to the reoptimization model?"

Bibliography

[AA+98] N. Alon, Y. Azar, G. J. Woeginger, and T. Yadid. "Approximation schemes for scheduling on parallel machines." *Journal of Scheduling* 1.1 (1998): 55-66.

[AE+14] F. N. Abu-Khzam, J. Egan, M. R. Fellows, Frances A. Rosamond, and Peter Shaw. "On the Parameterized Complexity of Dynamic Problems with Connectivity Constraints." *Combinatorial Optimization and Applications.* Springer International Publishing, 2014. 625-636.

[B88] H. L. Bodlaender. *Dynamic programming on graphs with bounded treewidth.* Springer Berlin Heidelberg, 1988.

[BF+13] M. A. Bender, M. Farach-Colton, S. Fekete, J. T. Fineman, and S. Gilbert. "Reallocation problems in scheduling." *In Proceedings of the twenty-fifth annual ACM symposium on Parallelism in algorithms and architectures*, pp. 271-279. ACM, 2013.

[BT14] G. Baram and T. Tamir. "Reoptimization of the minimum total flow-time scheduling problem." *Sustainable Computing: Informatics and Systems* 4.4 (2014): 241-251.

[CK+15] D. Chakrabarty, S. Khanna, and S. Li. "On $(1, \epsilon)$-restricted assignment makespan minimization." *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms.* SIAM, 2015.

[DMN12] J. Diaz, C. Munoz-Caro, and A. Nino. "A survey of parallel programming models and tools in the multi and many-core era." Parallel and Distributed Systems, IEEE Transactions on 23.8 (2012): 1369-1386.

[D59] Dijkstra, E. W. "A note on two problems in connexion with graphs." *Numerische mathematik* 1.1 (1959): 269-271.

[DF12] R. G. Downey and M. R. Fellows. *Parameterized complexity.* Springer Science & Business Media, 2012.

[EK72] J. Edmonds and R. M. Karp. "Theoretical improvements in algorithmic efficiency for network flow problems." *Journal of the ACM (JACM)* 19.2 (1972): 248-264.

[EKS08] T. Ebenlendr, M. Křćal, and J. Sgall. "Graph balancing: a special case of scheduling unrelated parallel machines." Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms. Society for Industrial and Applied Mathematics, 2008.

[GJ79] M. R. Garey and D. S. Johnson. "Computers and intractability: a guide to the theory of NP-completeness. 1979." *San Francisco, LA: Freeman* (1979).

[GL+04] M. Gairing, T. Läcking, M. Mavronicolas, and B. Monien. "Computing Nash equilibria for scheduling on restricted parallel links." *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. ACM, 2004.

[GMW07] M. Gairing, B. Monien, and A. Woclaw. "A faster combinatorial approximation algorithm for scheduling unrelated parallel machines." *Theoretical Computer Science* 380.1 (2007): 87-99.

[H35] P. Hall. "On representatives of subsets." *J. London Math.* Soc 10.1 (1935): 26-30.

[H96] D. S. Hochbaum. "Various notions of approximations: Good, better, best, and more." *Approximation algorithms for NP-hard problems*. PWS Publishing Co., 1996.

[HS76] E. Horowitz and S. Sahni. "Exact and approximate algorithms for scheduling nonidentical processors." *Journal of the ACM (JACM)* 23.2 (1976): 317-327.

[HS88] D. S. Hochbaum and D. B. Shmoys. "A polynomial approximation scheme for scheduling on uniform processors: Using the dual approximation approach." *SIAM journal on computing* 17.3 (1988): 539-551.

[J10] K. Jansen. "An EPTAS for scheduling jobs on uniform processors: using an MILP relaxation with a constant number of integral variables." *SIAM Journal on Discrete Mathematics* 24.2 (2010): 457-485.

[J15] J. Monnot. "A note on the traveling salesman reoptimization problem under vertex insertion." *Information Processing Letters* 115.3 (2015): 435-438.

[JLR15] K. Junosza-Szaniawski, M. Liedloff, and P. Rzażewski. "Fixing Improper Colorings of Graphs." *SOFSEM 2015: Theory and Practice of Computer Science*. Springer Berlin Heidelberg, 2015. 266-276.

[K84] N. Karmarkar. "A new polynomial-time algorithm for linear programming." Proceedings of the sixteenth annual ACM symposium on Theory of computing. ACM, 1984.

[LL+12] K. H. Lee, Y. J. Lee, H. Choi, Y. D. Chung, and B. Moon. "Parallel data processing with MapReduce: a survey." *AcM sIGMoD Record* 40.4 (2012): 11-20.

[LK79] J. K. Lenstra and A. R. Kan. "Computational complexity of discrete optimization problems." *Annals of Discrete Mathematics* 4 (1979): 121-140.

[LST90] J. K. Lenstra, D. B. Shmoys, and É. Tardos. "Approximation algorithms for scheduling unrelated parallel machines." *Mathematical programming* 46.1-3 (1990): 259-271.

[M11] D. Marx. "Fixed-parameter tractable scheduling problems." *Packing and Scheduling Algorithms for Information and Communication Services (Dagstuhl Seminar 11091).* Vol. 1. 2011.

[MV80] S. Micali and V. V. Vazirani. "An $O(\sqrt{|V|}|E|)$ algoithm for finding maximum matching in general graphs." Foundations of Computer Science, 1980., 21st Annual Symposium on. IEEE, 1980.

[MW13] M. Mnich and A. Wiese. "Scheduling Meets Fixed-Parameter Tractability." *arXiv preprint arXiv*:1311.4021 (2013).

[N06] R. Niedermeier. "Invitation to fixed-parameter algorithms." (2006).

[P14] D. R. Page. "Approximation Algorithms for Subclasses of the Makespan Problem on Unrelated Parallel Machines with Restricted Processing Times." *SOP Transactions on Applied Mathematics*, In Press, 2014.

[S12] O. Svensson. "Santa claus schedules jobs on unrelated machines." *SIAM Journal on Computing* 41.5 (2012): 1318-1341.

[ST93] D. B. Shmoys and É. Tardos. "An approximation algorithm for the generalized assignment problem." *Mathematical Programming* 62.1-3 (1993): 461-474.

[STT12] H. Shachnai, G. Tamir, and T. Tamir. "A theory and algorithms for combinatorial reoptimization." *LATIN 2012: Theoretical Informatics.* Springer Berlin Heidelberg, 2012. 618-630.

[SV05] E. V. Shchepin and N. Vakhania. "An optimal rounding gives a better approximation for scheduling unrelated machines." *Operations Research Letters* 33.2 (2005): 127-133.

[T77] K. Truemper. "On max flows with gains and pure min-cost flows." *SIAM Journal on Applied Mathematics* 32.2 (1977): 450-456.

[V13] V. V. Vazirani. *Approximation algorithms.* Springer Science & Business Media, 2013.

[VKW14] A. Verma, M. Korupolu, and J. Wilkes. "Evaluating job packing in warehouse-scale computing." *Cluster Computing (CLUSTER), 2014 IEEE International Conference on.* IEEE, 2014.

[VMH14] N. Vakhania, J. A. Hernandez, and F. Werner. "Scheduling unrelated machines with two types of jobs." *International Journal of Production Research* 52.13 (2014): 3793-3801.

[VP+15] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes. "Large-scale cluster management at Google with Borg." *In Proceedings of the Tenth European Conference on Computer Systems*, p. 18. ACM, 2015.

[VW14] J. Verschae and A. Wiese. "On the configuration-LP for scheduling on unrelated machines." *Journal of Scheduling* 17.4 (2014): 371-383.